Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty*

Hui Chen† Michael Michaux‡ Nikolai Roussanov§

August 7, 2018

Abstract

Mortgage refinancing activity associated with extraction of home equity contains a strongly counter-cyclical component consistent with household demand for liquidity. We estimate a structural model of liquidity management featuring counter-cyclical idiosyncratic labor income uncertainty, both long-term and short-term mortgages, and realistic borrowing constraints. We then empirically evaluate its predictions for the households' choices of leverage, liquid assets, and mortgage refinancing using micro-level data. Taking the observed historical paths of house prices, aggregate income, and interest rates as given, the model quantitatively accounts for many salient features in the evolution of balance sheets and consumption in the cross section of households over the 2001-2012 period.

JEL Codes: E21, E44, G21

Keywords: mortgage refinancing, home equity, housing collateral, liquidity constraints, household consumption and saving decisions, leverage

---

*We gratefully acknowledge comments and suggestions by Andy Abel, Rui Albuquerque, Fernando Alvarez, Nick Bloom, Paco Buera, John Campbell, Chris Carroll, Joao Cocco, Dean Corbae, Morris Davis, John Driscoll, Joao Gomes, Lars Hansen, Erik Hurst, Urban Jermann, Greg Kaplan, Ralph Koijen, Dirk Krueger, David Laibson, Francis Longstaff, Debbie Lucas, Hanno Lustig, Rajnish Mehra, Stijn Van Nieuwerburgh, Dimitris Papanikolaou, Jonathan Parker, Monika Piazzesi, Vincenzo Quadrini, Victor Rios-Rull, Tom Sargent, Martin Schneider, Antoinette Schoar, Amit Seru, Todd Sinai, Nick Souleles, Kjetil Storesletten, Harald Uhlig, Luis Viceira, Gianluca Violante, Jessica Wachter, Annette Vissing-Jorgensen, Pierre-Olivier Weill, Toni Whited, Paul Willen, Randy Wright, Amir Yaron, and audiences at a number of institutions and conferences.

†MIT and NBER

‡Orsus Research

§University of Pennsylvania and NBER

1 Introduction

The origins of the recent financial crisis and the severity of the Great Recession are often attributed to the increase in consumer indebtedness during the period of house price run-up in mid-2000s and the subsequent deterioration of household balance sheets with the sharp decline in house prices (see e.g., Dynan (2012), Mian, Rao, and Sufi (2013)). There is less consensus about the structural forces driving the borrowing boom and the subsequent consumption slump.[1] In particular, the expansion of household leverage and growth of consumer expenditures financed with extracted home equity over the period of house price boom as documented by Mian and Sufi (2010) are *qualitatively* consistent with liquidity-constrained households taking advantage of relaxed housing collateral constraints, but also with consumers' lack of self-control (e.g., as in Laibson (1997)), over-optimistic expectations (e.g., Laibson and Mollerstrom (2010)), and/or lender moral hazard (e.g., Keys, Mukherjee, Seru, and Vig (2010)).[2]

We document that mortgage refinancing activity involving home equity extraction exhibits a strongly counter-cyclical component that cannot be explained by fluctuations in interest rates, which suggests household demand for liquidity as an important driver of borrowing behavior. We show that a rational model of home equity-based borrowing by liquidity-constrained households that matches this key stylized fact can also *quantitatively* account for the empirical patterns in household leverage and consumption over the last decade.

In our model, households face idiosyncratic labor income risk and housing collateral constraints that resemble key institutional features of the U.S. mortgage markets. Specifically, households have access to long-term fixed rate mortgages and short-term home-equity lines of credit (HELOC), and they face two realistic borrowing constraints that restrict the ratios of loan size to home value (LTV) and loan size to household income (LTI) to be not too

[1] Keys, Piskorski, Seru, and Vig (2012) survey the extensive literature on the role of mortgage finance in the housing boom and bust.

[2] Landvoigt (2017) attributes the increase in homeowner leverage to rising uncertainty about future house prices rather than inflated growth expectations.

high at the time of new loan origination and refinancing. Additionally, our model features counter-cyclical idiosyncratic labor income risk (Meghir and Pistaferri (2004), Storesletten, Telmer, and Yaron (2004), Guvenen, Ozkan, and Song (2012)). This property of the labor income process implies that a macroeconomic downturn not only makes more households become liquidity constrained, but also increases their uncertainty about future income.

Our analysis focuses on households' optimal choices of consumption, leverage, precautionary savings in liquid assets and illiquid home equity, as well as the dynamic decisions in debt repayment, mortgage refinancing, home equity extraction, and default. We follow the partial equilibrium approach of Campbell and Cocco (2003).[3] While much of the existing literature treats mortgage refinancing and home-equity-backed borrowing in isolation, our analysis indicates that an integrated approach is important for understanding both.[4]

The decision to refinance trades off the benefits, in the form of lower interest rates and/or liquidity extraction, against the costs of originating a new loan, both financial and non-pecuniary. Because households do not have access to complete financial markets, the embedded options to default, prepay, or refinance the mortgage can no longer be analyzed in the standard option-pricing framework (see e.g., Chen, Miao, and Wang (2010)). As first pointed out by Hurst and Stafford (2004), the ability to convert home equity into liquid assets can generate refinancing even if it results in an increase in borrowing costs, which is in sharp contrast to the predictions of traditional models that consider lowering the interest rate as the only reason to refinance. Such liquidity-driven refinancing motives become particularly acute during a recession, when many households are subjected to large negative income shocks.

Fluctuations in household income and house prices also affect the tightness of households'

[3]We abstract from the choice between adjustable and fixed-rate mortgages analyzed by Campbell and Cocco (2003) and Koijen, Van Hemert, and Van Nieuwerburgh (2009). Our approach is also related to models of consumption smoothing in the presence of transaction costs, e.g. Bertola, Guiso, and Pistaferri (2005), Alvarez, Guiso, and Lippi (2010), and Kaplan and Violante (2011).

[4]The wealth and collateral effects of housing on consumption have been studied empirically (e.g. Caplin, Freeman, and Tracy (1997), Campbell and Cocco (2007), Carroll, Otsuka, and Slacalek (2011), Lustig and Van Nieuwerburgh (2010), Case, Quigley, and Shiller (2011), and Calomiris, Longhofer, and Miles (2012)), as well as theoretically (e.g., Campbell and Hercowitz (2005), Fernandez-Villaverde and Krueger (2011), Attanasio, Leicester, and Wakefield (2011), Favilukis, Ludvigson, and Van Nieuwerburgh (2011), and Midrigan and Philippon (2011)).

borrowing constraints. A rise in house prices relaxes the LTV constraint, resulting in an increase in cash-outs for high-leverage households. A rise in household income, on the other hand, relaxes the LTI constraint, enabling more low-income households to access their home equity savings. Moreover, a looser LTI constraint can also enable more households to become homeowners or switch to a bigger house, which in turn relaxes the LTV constraint and further increases the amount of borrowing.[5] Thus our model highlights household demand as a key force behind the strong credit expansion during the house price boom from 2000 to 2006.

Taking the observed historical paths of house prices, aggregate household income, and interest rates as exogenously given, our baseline model can account for both the run-up in household leverage from 2000 to 2006, and the sharp contraction in consumption that followed. In the cross section, absent any ex ante heterogeneity, the model generates wide dispersion in the dynamics of liquid asset holding, household debt, refinancing patterns, as well as consumption that are largely consistent with the data.

Specifically, the counter-cyclical idiosyncratic labor income risks, borrowing constraints that vary in tightness with house prices and income, and high costs of default together generate a strong precautionary saving motive, especially for high-leverage households. As a result, in spite of the buffer provided by long-term mortgages and HELOCs, households with high boom-time leverage experience significantly larger consumption declines and debt reductions during the Great Recession in our model. In other words, deleveraging need not be "forced" as in the case of short-term mortgages (contrary to implications of more stylized models, e.g. Justiniano, Primiceri, and Tambalotti (2013)). Furthermore, in contrast to the findings of Kaplan, Mitman, and Violante (2017), the dispersion in consumption responses in our model is relatively small across households that differ in the size of housing shocks (relative to total wealth) they experience. Thus, in our model the drop in household consumption around the housing crisis is not just due to a loss of housing wealth; the interaction between

[5] Campbell and Cocco (2015) emphasize the role of the LTI ratio in driving borrower default decisions. Greenwald (2017) studies the effect of the interactions between LTV and debt service constraints on housing demand and house prices in a general equilibrium setting.

the house price shocks and the financing frictions associated with household debt plays an important role. These results, while in partial equilibrium, suggest that the effect of the subsequent macroeconomic slowdown is amplified by greater indebtedness accumulated during the boom years.

To discipline the model, we estimate its structural parameters by targeting the key moments of household consumption, asset and debt holdings, as well as the aggregate dynamics of mortgage refinancing and equity extraction in relation to macroeconomic conditions. We show that, in the presence of borrowing constraints, including the counter-cyclical dynamics of cash-out refinancing among the target moments is key for identifying the parameters responsible for the strength of consumption smoothing motive, such as the intertemporal elasticity of substitution (IES). Borrowing constraints imply that household consumption is not very sensitive to interest rates. In fact, we estimate an IES that is well below unity (around 0.3) but much higher than that implied by the reduced-form sensitivity of aggregate consumption to interest rates, which is close to zero both in the model and in the data.

While we rely on aggregate cross-sectional and time series moments as targets in our structural estimation, we utilize a range of cross-sectional features constructed from the micro data to evaluate the model's predictions and learn about its limitations. Overall, our model does a good job reproducing the cross-sectional distribution of leverage (as measured by both LTV and LTI), although it understates its right tail by missing the extremely highly-levered households. The model implies that liquidity-driven refinancing behavior is especially prevalent among the constrained households. We show that this prediction is broadly consistent with the cross-sectional data. Our baseline model also predicts very few defaults. This is related to its failure to generate households with very high leverage, which is in part due to the tight LTI constraints that we impose, but also because our structural parameter estimates imply high costs of default perceived by households. Importantly, these high costs of default are endogenously generated, as the collateral value of housing raises the value of homeownership.[6]

Since our baseline model sets the LTV and LTI limits according to conforming mortgage lending standards, it likely provides a lower bound for the effects of income and house price shocks on household leverage expansion. This is because the loosening of mortgage lending standards in the early 2000s, e.g. via expansion of subprime and low-documentation loans, implies more marginal households would see their constraints relaxed than our baseline model allows (e.g., as in Favilukis, Ludvigson, and Van Nieuwerburgh (2011)). Similarly, the subsequent consumption drop could be even more drastic if lending standard were tightened (e.g., Guerrieri and Lorenzoni (2011), Midrigan and Philippon (2011)).[7] Indeed, we show that the fit with the data further improves when we relax the LTV and LTI constraints.[8] Similarly, our model shows that relaxing lending standards is necessary to explain the rise in foreclosure following the crash (e.g., Corbae and Quintin (2013)).

2 Empirical Evidence

In this section we document some new stylized facts on how households access liquidity in response to aggregate economic conditions via mortgage refinancing and other forms of home equity-based borrowing. The evidence not only motivates our modeling choice in Section

[6]While the structure of our model is similar to Campbell and Cocco (2015), we set tighter LTI and LTV constraints, following the GSE guidelines, and estimate the structural parameters by targeting a range of moments, which do not include default rates, given the different focus of our paper. Chatterjee and Eyigungor (2015) match both homeownership and mortgage default rates in a model with long-term loans but without the possibility of refinancing, thus understating the collateral value of housing and, consequently, the cost of default. Laufer (2017) matches default rates closely by using a model with regional and idiosyncratic house price shocks (as well as housing preference shocks and expectations shocks), all of which our model abstracts from, and estimated to fit the individual leverage ratios at origination that often exceed the constraints that we impose, while not allowing for interest rate fluctuations that are important in our setting.

[7]Carroll, Slacálek, and Sommer (2012) argue that an increase in labor income uncertainty, rather than the tightening of credit constraints by themselves, was the main driver of the consumption decline during the Great Recession.

[8]While our model takes the evolution of house prices as given, a number of authors have attributed part of the house price run-up to the easing of lending standards (e.g., Landvoigt, Piazzesi, and Schneider (2012)), and some of the subsequent crash to an exogenous tightening of credit (e.g., Favilukis, Ludvigson, and Van Nieuwerburgh (2011) and Midrigan and Philippon (2011)). In contrast, Kaplan, Mitman, and Violante (2017) argue that future housing demand expectations played by far the dominant role. See also Rios-Rull and Sanchez-Marcos (2008), Ortalo-Magné and Rady (2006), He, Wright, and Zhu (2012) for analyses of the endogenous evolution of house prices with collateral constraints.

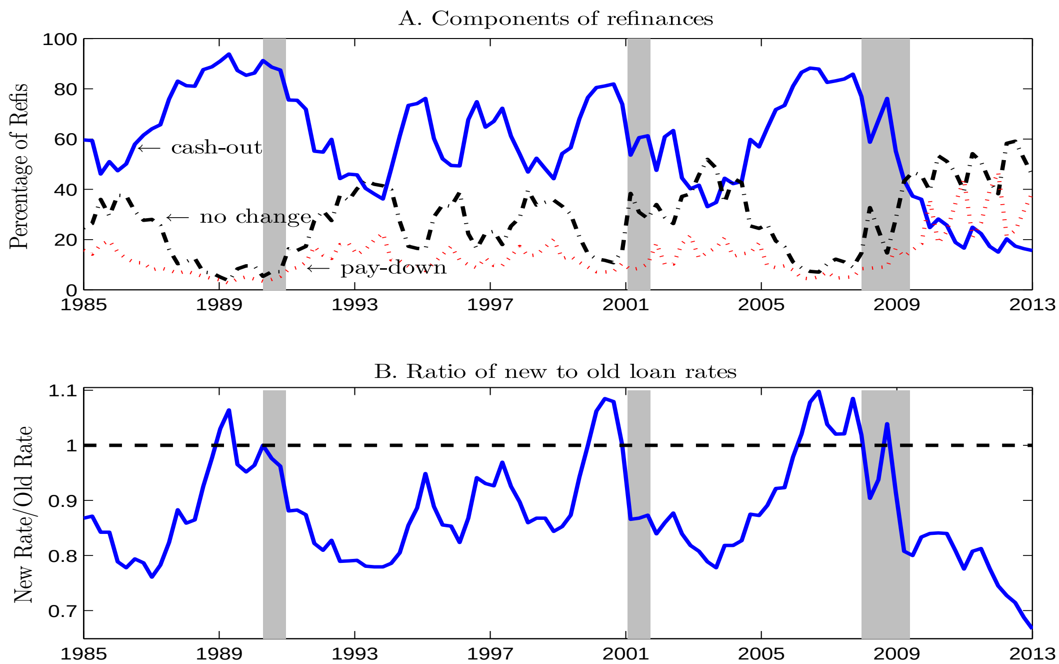

Figure 1: **Fraction of cash-out and the median rate ratio for refinance loans.** Panel A plots the percentage of refinancing resulting in 5% higher loan amount (cash-out), no change, or lower loan amount (pay-down). Panel B plots the median ratio of new to old loan rates upon refinance. The data is from Freddie Mac for the period 1985Q1 to 2012Q4.

3, but also helps identify the structural parameters when we take our model to the data in Section 4.

We begin by analyzing how mortgage refinancing activity at the aggregate level relates to interest rates and macroeconomic conditions. The refinancing measures we use are refinancing applications index from the Mortgage Bankers Association (the MBA refi index) and data on mortgage refinancing volume from Freddie Mac.

Figure 1 Panel A plots the time series of the percentage of originated refinance loans in the Freddie Mac data for which the loan amount (i) is raised by 5% or more (classified as "cash-out"), (ii) remains within 5% of the original amount (classified as "no change"), or (iii) is reduced by 5% or more (classified as "pay-down"). On average, 61% of refinancing over the period of 1985 to 2013 are cash-outs. The share of cash-outs is visibly higher towards the end of each economic expansion, and it declines coming out of recessions. In contrast, the fraction of refinancing that results in no change in loan balance or pay-down typically rises after the end of recessions, presumably because households refinance at such times to take advantage of the lower mortgage rates and to repay the loans they take out entering the recession.

Since the standard theory predicts that the primary driver of mortgage refinancing is to lower the borrowing costs, it is informative to examine under what conditions refinancing actually lowers borrowers' loan rates. Panel B of Figure 1 plots the median of the ratio of new mortgage rates to the old rates on the refinanced loans (adjustable rate mortgages are excluded). While households do refinance at lower rates coming out of a recession, they tend to refinance despite higher rates towards the end of an expansion and beginning of a recession, when the median rate ratio can sometimes exceed unity.

The evidence above clearly shows that interest savings are not the only driver of refinancing. Borrowers may try to alleviate liquidity constraints either by increasing the loan amount (cash-out) or extending the loan term (thus reducing the monthly payments).[9] Indeed, the correlation between the rate ratio and the cash-out share in Figure 1 is 78%. Given that labor income is not tradable and other uncollateralized personal loans (e.g., credit card loans) are expensive, mortgage loans are a major source of credit for liquidity constrained households. The finding is consistent with the household-level evidence in Hurst and Stafford (2004) that the most liquidity-constrained households refinance following negative income shocks even as interest rates increase.

Aggregate refinancing activity. To further investigate the dynamics of the aggregate refinancing activity, we regress the monthly MBA refi index on a host of financial and macroeconomic variables:

$$REFI_t = \beta_0^{REFI} + \beta_Z^{REFI} \Delta IP_t + \beta_H^{REFI} \Delta HPI_t + \beta_R^{REFI} R_t^{M30} + \beta_{\Delta R}^{REFI} \Delta R_t^{M30} + \beta_r^{REFI} r_t^{1Y} + \epsilon_t, \quad (1)$$

[9] Households are strictly worse off by refinancing into a higher rate loan in the case of "no-change" or "pay-down" refinancing as long as the loan's time to maturity remains the same. In the case of "pay-down", the households will be better off by choosing to prepay instead.

Table 1: Aggregate Refinancing and Home Equity Extraction

	A. Refinancing			B. Home Equity Extraction					
				Prime, first-lien mortgage			HEL+HELOC		
$\Delta macro$	-0.42***	-0.25***	-0.20**	-0.00	-0.12***	-0.13***	0.06	-0.01	-0.03
	(0.16)	(0.09)	(0.10)	(0.05)	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)
ΔHPI_t		0.15	0.16*		0.06***	0.06***		0.06***	0.06***
		(0.10)	(0.10)		(0.02)	(0.02)		(0.02)	(0.02)
R_t^{M30}		-1.91***	-1.98***		-0.43***	-0.43***		-0.04	-0.04
		(0.67)	(0.68)		(0.15)	(0.13)		(0.11)	(0.10)
ΔR_t^{M30}			-1.46*			0.21***			0.19***
			(0.85)			(0.08)			(0.06)
r_t^{1Y}		-1.16*	-0.99*		0.28***	0.26***		0.05	0.03
		(0.61)	(0.57)		(0.10)	(0.09)		(0.08)	(0.07)
$\bar{R}^2$	0.06	0.65	0.67	-0.06	0.49	0.55	0.11	0.61	0.68

NOTE: Results in Panel A are based on monthly data, January 1990 to December 2012. Numbers in parentheses are Newey-West standard errors with 12 lags. Results in Panel B are based on annual data, 1993 – 2012. Numbers in parentheses are Newey-West standard errors with 4 lags. Refinancing activities are measured by the scaled MBA refi index. Home equity extraction is measured by the ratio of annual dollar amount of cash-out from prime, first-lien conventional mortgages or home-equity loans and lines of credit (HEL+HELOC) to previous-year personal income. In Panel A (B), $\Delta macro$ is measured by the 12-month growth rate in industrial production (one-year real personal income growth), and ΔHPI_t is the one-year growth in the Case-Shiller house price index (FHFA house price index). R^{M30} and ΔR^{M30} are the 30-year conventional mortgage rate and its annual change. r^{1Y} is the 1-year constant maturity treasury yield.

where ΔIP_t is the year-on-year growth in Industrial Production, ΔHPI_t is the year-on-year growth in the Case-Shiller house price index, R_t^{M30} is the 30-year fixed mortgage rate, ΔR_t^{M30} is the year-on-year change in R_t^{M30}, and r_t^{1Y} is the 1-year Treasury rate. To make the coefficients easier to interpret, we rescale the MBA refi index to have a mean of 8%, the average annual refinancing rate for homeowners according to the Home Mortgage Disclosure Act (HMDA) data.

Panel A of Table 1 reports the results. Among the key drivers of mortgage refinancing are the current 30-year mortgage rate and its yearly change, both of which come in with negative and significant coefficients. This is consistent with households refinancing to take advantage of lower interest rates. Recent house price growth affects refinancing positively, since an increase in house prices implies an increase in home equity that can be cashed out, but the significance is marginal. The industrial production growth (represented by $\Delta macro$ in Panel A), a direct measure of economic activity, has a robustly significant and negative coefficient even after controlling for mortgage rates and house price growth. Thus, households refinance more in economic downturns, beyond what can be explained by the changes in interest rates.

One potential concern is that the measures of macroeconomic conditions could be a proxy for the effects of interest rates not captured by the term structure variables included in the regressions. In Appendix A, we present further evidence on counter-cyclical refinancing activity that exploits state-level variation in economic conditions and house prices, which is less synchronized with the fluctuations in interest rates at the national level.

Aggregate home equity withdrawal. The aggregate refinancing rate includes both loans that involve cash-out and those that do not. Next, we examine how the actual withdrawal of home equity by households relates to different macroeconomic conditions.

We separately examine two mechanisms of home equity withdrawal to differentiate the roles that senior and junior mortgage loans play in smoothing income shocks. Our first measure is cash-out refinancing of first-lien conforming mortgages, while the second combines home equity loans and lines of credit (HEL+HELOC, computed as the net change of the outstanding balances reported in the Flow of Funds). We normalize the dollar amount of total home equity withdrawn each year by the total personal income in the previous year and then regress it on real personal income growth, house price growth, and several interest rate variables as in the regression of refi rates:

$$HEW_t^j = \beta_0^j + \beta_Z^j \Delta PI_t + \beta_H^j \Delta HPI_t + \beta_R^j R_t^{M30} + \beta_{Rl}^j \Delta R_t^{M30} + \beta_r^j r_t^{1Y} + \epsilon_t. \tag{2}$$

where $j \in \{\text{Cash-out}, \text{HELOC}\}$, HEW_t is the home equity withdrawal in a year (via cash-out or HELOC) scaled by the total personal income in the previous year, ΔPI_t is the one-year growth rate in real personal income, $\Delta R_t^{M30} = R_t^{M30} - R_{t-1}^{M30}$, and the other variables are the same as defined in (1).

The results are shown in Panel B of Table 1. Like refinancing rates, cash-out volume is negatively related to the level of 30-year mortgage rate. However, it is positively related to the change in 30-year mortgage rate, the opposite of the case for refinancing (see Panel A). When households decide when to cash out, they not only compare the level of current mortgage rate to old rates, but also to the costs of other sources of financing (e.g., rates on credit card debt). Moreover, the degree of liquidity constraint households face is a key factor. The fact that cash-out refinancing occurs despite rising costs of mortgage borrowing is consistent with the observation that both the cash-out share and the rate ratio tend to peak at the ends of expansions and early in the recessions (see Figure 1). Finally, similarly to total refinancing, house price growth is positively related to both measures of home equity withdrawal, with an effect of essentially identical magnitude, indicating that out of an extra $1 of home equity 6 cents are withdrawn in the same year.

After controlling for house price growth and interest rates, growth in real personal income is significantly negatively correlated with cash-out from first-lien mortgages. If real income drops by 10%, households on average increase cash-out from their first-lien mortgages by about 1.3% of income to offset the drop. These aggregate sensitivities mask substantial underlying heterogeneity across households in terms of homeownership, income, leverage, and liquid assets, which potentially lead to different cash-out responses to income shocks. Such heterogeneous responses are the focus of our structural model.

Interestingly, equity withdrawal via home equity loans and lines of credit has no significant relation with aggregate personal income growth. This is in sharp contrast to the similar sensitivity that both measures of home equity withdrawal have for house price growth. This evidence suggests that households primarily use refinancing of senior-lien mortgage loans for consumption smoothing. This may be due to the fact that HEL(OC)s are relatively expensive due to their higher credit risk (they are junior and not insured by the GSEs) and shorter maturity (which generates rollover risk).[10] We will also use our structural model to

[10] In fact, households often use funds extracted upon cash-out to repay outstanding junior loans, as well as other forms of debt, such as credit card balances, as can be seen in the Survey of Consumer Finances (SCF).

understand the distinctive properties of the two methods of home equity extraction.

3 Quantitative Model

In this section, we present a dynamic model of household consumption, saving, and borrowing decisions with incomplete markets. Households are confronted with idiosyncratic shocks to income and aggregate shocks to interest rates, income growth, and house value. Since our focus is to capture households' behavior in the face of realistic macroeconomic risks and constraints, we try to model the key elements of the institutional environment of the U.S. housing finance while taking asset prices (including house prices) as exogenous.

3.1 Model specification

The economy is populated by ex-ante identical, infinitely lived households, indexed by i. We assume households have recursive utility over real consumption as in Epstein and Zin (1989) and Weil (1990),

$$U_{i,t} = \left[(1-\delta) X_{i,t}^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E}_t \left[U_{i,t+1}^{1-\gamma} \right]^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \tag{3}$$

where δ is the time discount rate, γ is the coefficient of relative risk aversion, ψ is the intertemporal elasticity of substitution (IES), $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$, and $X_{i,t}$ is a Cobb-Douglas aggregator of housing services $s_{i,t}$ and real non-housing consumption $c_{i,t}$,[11]

$$X_{i,t} = s_{i,t}^{\nu} c_{i,t}^{1-\nu}.$$

In the special case with $\theta = 1$, we recover CRRA utility.

The nominal price level at time t is P_t. We assume that the (gross) inflation rate is constant, $P_{t+1}/P_t \equiv \pi$. It is important to model inflation since mortgage loan contracts are

[11] Piazzesi, Schneider, and Tuzel (2007) argue for a preference structure that is close to Cobb-Douglas based on the joint behavior of the U.S. housing expenditure shares and asset prices over time, while Davis and Ortalo-Magne (2011) show that a Cobb-Douglas specification is broadly consistent with the cross-sectional U.S. data.

nominal and therefore their balances are eroded over time by inflation, increasing home equity even in the absence of principal amortization. At the same time, inflation variability over the time period that we focus on in our estimation is relatively mild, allowing us to abstract from inflation risk in order to contain the dimensionality of the state space.

Each household is endowed with one unit of labor supplied inelastically, which generates before-tax nominal income $y_{i,t}$. The income tax rate is τ. We assume $y_{i,t}$ has an aggregate real income component, Y_t, an idiosyncratic component, $\tilde{y}_{i,t}$, as well as adjustment for inflation:

$$y_{i,t} = P_t \; Y_t \; \tilde{y}_{i,t}. \tag{4}$$

The growth rate of aggregate real income is $Z_{t+1} = Y_{t+1}/Y_t$. The idiosyncratic labor income component, $\tilde{y}_{i,t}$, follows an autoregressive process with state-dependent conditional volatility,

$$\log \tilde{y}_{i,t} = \log \mu_y(Z_t) + \rho_y \log \tilde{y}_{i,t-1} + \sigma(Z_t)\epsilon^y_{i,t}, \quad \epsilon^y_{i,t} \sim \mathcal{N}(0,1). \tag{5}$$

The counter-cyclical nature of idiosyncratic labor income risk, which is captured here by having $\sigma(Z_t)$ decreasing in Z_t, is emphasized by Storesletten, Telmer, and Yaron (2004). We set $\log \mu_y(Z) = -\frac{1}{2}\frac{\sigma^2(Z)}{1+\rho_y}$, so that the cross-sectional mean of $\tilde{y}_{i,t}$ is normalized to 1.

Next, we specify households' assets, liabilities, and the financing constraints.

Liquid assets Households have access to a riskless savings account with balance $a_{i,t}$, which earns the nominal short rate r_t. Interest income is taxed at the same rate τ as labor income. We also refer to the savings account as the households' liquid assets, in contrast to the illiquid housing assets.

Houses A household can choose to own $h_{i,t}$ units of housing, with $h_{i,t} \in \{h_1, \cdots, h_n\}$, which generates housing service flow $s_{i,t} = h_{i,t}Y_t$. Indexing per-unit housing service to real aggregate income Y_t ensures that aggregate housing and non-housing consumption are consistent with balanced growth.

Houses are valued proportionally at price P_t^H per unit. We assume that the nominal house price level P_t^H is co-integrated with the nominal aggregate income, $P_t Y_t$. Specifically,

$$P_t^H = \bar{H}\ P_t\ Y_t\ p_t^H, \tag{6}$$

where $\bar{H}$ is the long-run house price-to-income ratio, while p^H is a stationary process that represents the aggregate risk inherent in the housing market's transitory deviations from the trend in aggregate income. Finally, the sale or purchase of a home incurs a proportional transaction cost ϕ_h.[12]

Debt There are two types of borrowing allowed for households, both of which are collateralized by the house: long-term fixed-rate mortgages and short-term home equity lines of credit (HELOC). We assume that long-term mortgage contracts are perpetual interest-only mortgages. This assumption is reasonable since households are infinitely lived in our model, and implies that a household could potentially maintain a constant nominal level of mortgage debt over time, only paying interest on its borrowing. This assumption could potentially understate the amount of home equity extraction relative to the data since households are not forced to accumulate equity in the first place (other than through the inflation channel discussed above). The coupon rate for mortgages originated in period t is R_t, which can be different from the coupon rate for existing mortgages, $k_{i,t}$.[13]. Based on the beginning-of-period mortgage balance $b_{i,t}$ and coupon rate $k_{i,t}$, the mortgage payment in period t is $k_{i,t}b_{i,t}$. Households can deduct the mortgage interest expense, which is the full mortgage payment for an interest-only mortgage, from their taxable income $y_{i,t}$.

[12]Our approach implicitly treats house size as fundamentally limited by the availability of fixed factors such as land, similarly to the approaches in Ortalo-Magné and Rady (2006) and Corbae and Quintin (2013). Alternatively, one can model housing stock as fully adjustable through investment and depreciation, e.g. as in Favilukis, Ludvigson, and Van Nieuwerburgh (2011) and Iacoviello and Pavan (2013). Kiyotaki, Michaelides, and Nikolov (2011) consider the combination of both fixed and adjustable factors in the total value of the housing stock.

[13]We do assume that this rate is faced by all households trying to borrow at time t. This is broadly consistent with the fact that government-backed conforming mortgage rates exhibited very little variation, as documented, e.g., in Hurst, Keys, Seru, and Vavra (2014).

The HELOC is modeled as a one-period debt with floating interest rate benchmarked to the riskfree rate r_t, $r_t^{HL} = r_t + \vartheta$, with spread $\vartheta > 0$ over the short rate r_t. It is costless to adjust the HELOC balance, although the balance is subject to a set of borrowing constraints every period, which we specify below. Due to the interest rate spread ϑ and the borrowing constraints, it is never optimal to simultaneously hold non-zero balances in HELOC and liquid assets. Thus, we can capture the HELOC and liquid asset balance with the same variable $a_{i,t}$. Specifically, the balance of HELOC and liquid assets are $-a_{i,t}^-$ and $a_{i,t}^+$, respectively, with $a_{i,t}^+ = \max(a_{i,t}, 0)$ and $a_{i,t}^- = \min(a_{i,t}, 0)$.

When a homeowner sells the home and become a renter, it immediately repays all the outstanding debt – including the current period mortgage coupon payment, the remaining mortgage balance, and the HELOC balance – using the net proceeds of house sale and its stock of liquid assets.

Mortgage refinancing and repayment Households have the option to refinance the long-term mortgage, which results in a reset of the coupon rate $k_{i,t+1}$ from $k_{i,t}$ to the current market mortgage rate R_t, as well as a possibly different mortgage balance $b_{i,t+1}$. In particular, a cash-out refinancing is one that results in a higher mortgage balance, $b_{i,t+1} > b_{i,t}$.

When a household refinances into a new loan with balance $b_{i,t+1}$, they will incur a cost equal to $\phi(b_{i,t+1}; S_t)$. Therefore, the net proceeds from refinancing will be $b_{i,t+1} - b_{i,t} - \phi(b_{i,t+1}; S_t)$. The refinancing costs include the opportunity cost of time spent on the refinancing process, which does not depend on the loan amount, as well as direct fees associated with issuing a new mortgage, which tend to scale with the loan size. The cost of refinancing has both a quasi-fixed component (indexed to nominal aggregate income) and a proportional component:

$$\phi(b_{i,t+1}; S_t) = \phi_0 \; P_t Y_t + \phi_1 \; b_{i,t+1}. \tag{7}$$

Besides refinancing, households can also reduce their mortgage balance costless ly at any time by repaying the mortgage, i.e., choosing $b_{i,t+1} < b_{i,t}$, which does not change the existing coupon rate, $k_{i,t+1} = k_{i,t}$.

Collateral and debt service constraints When households apply for new loans, they face a pair of borrowing constraints: the *loan-to-value constraint* (LTV) and the *loan-to-income constraint* (LTI). Specifically, these constraints are imposed when the new HELOC balance is non-zero ($a^-_{i,t+1} < 0$), or when the household obtains a new mortgage, which occurs when they buy a new house or refinance the existing mortgage.

The LTV constraint restricts the new combined balances of all loans, including mortgage and HELOC, relative to the house value:

$$b_{i,t+1} - a^-_{i,t+1} \leq \xi_{LTV}\ P^H_t h_{i,t}, \tag{8}$$

with $\xi_{LTV} \geq 0$. Similarly, the LTI constraint restricts the new combined balances of all loans relative to household nominal income:

$$b_{i,t+1} - a^-_{i,t+1} \leq \xi_{LTI}\ y_{i,t}, \tag{9}$$

with $\xi_{LTI} \geq 0$. The constraints (8) and (9) mimic the loan-to-value and debt-to-income constraints widely used in practice, in particular, for conforming loans.

In addition, we impose an upper bound on the HELOC balance (or a lower bound on $a_{i,t}$) as a fraction $-\underline{a}$ of permanent income,

$$- a^-_{i,t+1} \leq -\underline{a} P_t Y_t. \tag{10}$$

This constraint is motivated by the common practice that limits the size of HELOCs and home equity loans to reduce the risk of default.

Default Homeowners have the option to default on their mortgages and HELOCs. When a household defaults on any of its debt, its home is ceased and it becomes a renter. Furthermore, the defaulted household will be excluded from the housing market for a stochastic period of time. With probability ω each period, it will regain eligibility for becoming a homeowner, at which point the household can choose to buy a house or remain a renter. This approach of modeling homeownership and default decision broadly follows Campbell and Cocco (2015).

Renting Unlike homeowners, a renter household can freely adjust the amount of housing services it consumes each period. For simplicity, we assume the ratio of rent per unit of housing relative to nominal aggregate income is a constant ϖ. The parameter ϖ can also capture the disutility of renting relative to owning a home. An unrestricted renter (not excluded from the housing market due to default) can become a homeowner by purchasing a house, using savings and borrowing.

3.2 Summary of exogenous shocks

In total, there are three aggregate state variables, summarized in the aggregate state vector $V_t = (Z_t, p_t^H, r_t)$. We assume that V_t follows a first-order vector autoregressive process (VAR) in logarithms:

$$\log V_{t+1} = \mu_V + \Phi_V \log V_t + \sqrt{\Sigma_V} \epsilon_{t+1}^V. \tag{11}$$

We assume that the mortgage rate R_t is a function of the aggregate state variables. We choose the following linear-quadratic specification for R_t, which is motivated empirically (see Section 4.1):

$$\log R(V_t) = \kappa_0 + \kappa_1' \log V_t + \kappa_2 \left(\log p_t^H\right)^2. \tag{12}$$

For an individual household, the vector of exogenous state variables, denoted by $v_{i,t}$, contains the individual labor income and the aggregate state vector: $v_{i,t} \equiv (y_{i,t}, V_t)$.[14]

We characterize the intertemporal optimization problem for homeowners and renters using standard dynamic programming tools, as detailed in Appendix B.

[14] We assume that all households bear the same aggregate risks since we focus on the "average" household that is likely to need to use home equity to smooth consumption. There is some evidence in the recent literature that wealthier households are disproportionately affected by aggregate fluctuations, see e.g., Parker and Vissing-Jørgensen (2009).

4 Structural Estimation

This section describes the empirical implementation of the model in Section 3. To solve the model, we discretize the state space and apply standard numerical dynamic programming techniques. We estimate the model parameters in three steps. First, we specify the dynamics of the exogenous state variables based on empirical estimates. Second, we set the institutional parameters to broadly represent the environment faced by U.S. households. Third, we estimate the preference and transaction cost parameters by matching the model-implied moments (computed from the simulation of a large panel of households) of household assets, liabilities, and consumption, as well as the dynamics of mortgage refinancing, with the data, taking the pre-estimated state variable dynamics and pre-set institutional parameters as given. Thus, our approach is essentially a version of the simulated method of moments (e.g., Duffie and Singleton (1993)) where a set of "nuisance" parameters are pre-specified before the structural parameters are estimated.[15] Details of the procedure can be found in Appendix C.

4.1 Exogenously specified parameters

Aggregate state variable dynamics We first estimate the VAR for the aggregate state variables in (11) using annual data. To reduce the degrees of freedom, we impose the restriction that Φ_V is diagonal. We use the U.S. real GDP growth rate as proxy for the real growth rate in aggregate income Z_t in the model, the one-year Treasury bill rate as proxy for the nominal short rate r_t, and the demeaned log house price-GDP ratio (computed using the S&P Case-Shiller house price index and GDP data) as proxy for the transitory component in house price h_t. The estimated parameters of the VAR are reported in Table 2. We then approximate the VAR with a discrete-state Markov chain using the method of Tauchen and

[15] Dridi, Guay, and Renault (2007) provide a formal justification of this approach based on the indirect inference methodology (Smith (1993), Gallant and Tauchen (1996), and Gourieroux, Monfort, and Renault (1993)). Laibson, Repetto, and Tobacman (2007) follow a similar strategy for estimating the structural parameters in a household consumption and liquidity management model with hyperbolic discounting. Gourinchas and Parker (2002) pioneered structural estimation of household consumption-saving models. Hennessy and Whited (2005) apply structural estimation in corporate debt and investment models.

Table 2: Aggregate State Variables

Panel A: VAR Parameters

	μ	Φ_s			$\Sigma_s \times 10^{-3}$		
GDP	0.013	0.420	0	0	0.492	0.576	0.006
p_t^H	-0.015	0	0.888	0	0.576	6.525	0.440
r_t	0.002	0	0	0.844	0.006	0.440	0.192

Panel B: Mortgage Rate Parameters

κ_0	κ_Z	κ_{p^H}	κ_r	$\kappa_{(p^H)^2}$	R^2
0.049	0.094	0.011	0.684	-0.270	0.949
(0.001)	(0.023)	(0.004)	(0.025)	(0.022)	

Hussey (1991). The state variables (Z, p^H, r) are discretized using 2, 10, and 10 grid points, respectively.

Panels A-C of Figure 2 compares the actual time series of the three aggregate state variables (blue solid lines) against the Markov chain approximation (red circle lines) for the period 1987-2012. Panel A shows that the 2-state approximation tracks the history of real income growth well over all, but it understates the severity of the Great Recession and slightly overstates the extent of the recovery thereafter. Panel B and C show that our model captures closely the highly persistent deviations of house prices from the trend of real economic growth and the paths of nominal short-term rates.

For tractability, we specify the mortgage rate R_t as an exogenous quadratic function of all the aggregate state variables as in Equation (12). Panel C of Table 2 reports the regression estimates of this relation based on the 30-year conforming mortgage rate (our empirical proxy for R). We obtain an R^2 of 95% with just 4 explanatory variables $(Z_t, p_t^H, r_t, (p_t^H)^2)$, suggesting that this exogenous function $R(V)$ captures most of the time variation in the long-term mortgage rate. Since the household's fixed mortgage rate $k_{i,t}$ is part of the endogenous state variables that spans the same states as R_t, in order to keep the size of the state space manageable we use a coarser grid for the latter with 7 points based on the implied

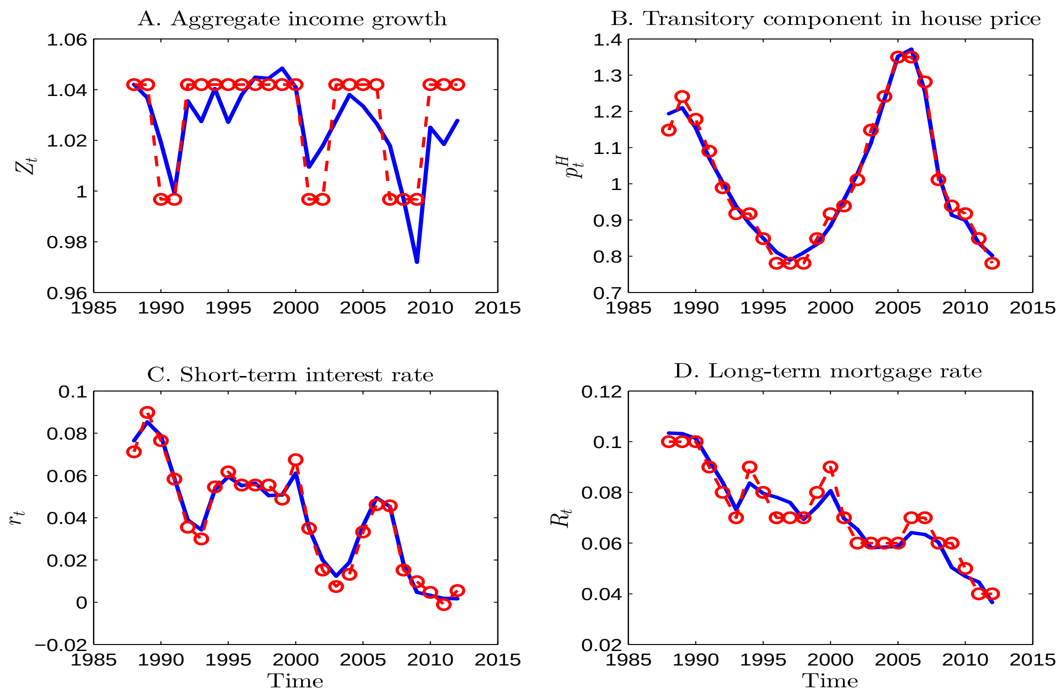

Figure 2: **Time series of exogenous state variables.** Solid lines: data; circles indicate corresponding approximated values in the model.

distribution of $R(V)$. Panel D of Figure 2 plots the long-term mortgage rate in the data and the corresponding value on the grid. The discretized process for R_t tracks the history of the mortgage rates closely throughout the sample.

We set the set of possible house sizes that a household can choose from to $\{0.75, 1, 1.25\}$. The choice of $\bar{H} = 4$ is based on estimates obtained using micro data (in the Survey of Consumer Finances for 2001, a year when the house price to GDP ratio is close to its long-run mean, the average ratio of housing assets to income among homeowners with positive income equals approximately 3.95). Finally, given the relatively smooth evolution of inflation over the sample period, we assume a constant inflation rate equal to its historical average $\pi = 2.85\%$ per annum.

Idiosyncratic state variable dynamics We calibrate the process for the idiosyncratic component of labor income $\tilde{y}_{i,t}$ (5) following Storesletten, Telmer, and Yaron (2007). With two states for the growth rate of real aggregate income, we set the conditional volatility of

Table 3: Parameter Values

This table reports the parameter values for the model. For the estimated parameters, the values in parentheses are the standard errors.

Panel A. Exogenously-fixed parameters

ρ_y	$\sigma(Z_G)$	$\sigma(Z_B)$	τ	$\bar{H}$	ξ_{LTV}	ξ_{LTI}	$-\underline{a}$	ω	ζ	ϑ
0.95	0.12	0.21	0.25	4.00	0.80	3.50	0.30	0.15	1.0	0.04

Panel B. Estimated parameters

δ	γ	ψ	ν	ϖ	ϕ_0	ϕ_1	ϕ_h
0.920	3.036	0.301	0.134	1.324	0.154	0.014	0.135
(0.007)	(0.347)	(0.020)	(0.004)	(0.100)	(0.020)	(0.008)	(0.017)

$\tilde{y}_{i,t}$ to $\sigma(Z_G) = 12\%$ and $\sigma(Z_B) = 21\%$. The autocorrelation parameter is $\rho_y = 0.95$. This process is then discretized as a Markov chain with 12 grid points.

Institutional parameters Several exogenously set parameters reflect the main institutional features of the U.S. economy for homeowners and renters. The personal income tax rate is $\tau = 25\%$. The set of borrowing constraints includes (i) the constraint on the loan-to-value ratio $\xi_{LTV} = 80\%$, (ii) the constraint on the loan-to-income ratio $\xi_{LTI} = 3.5$, both of which are broadly consistent with the conforming loan requirements, and (iii) the upper bound on HELOC balances is $-\underline{a} = 30\%$ of aggregate income. The period of exclusion from debt markets for defaulted households is on average 7 years, as represented by the annual probability of $\omega = 0.15$ for returning to the housing market. Finally, we set $\zeta = 1$, so that a household does not lose any of its liquid assets at default. Most of these parameter choices closely follow Campbell and Cocco (2015).

The idiosyncratic labor income and institutional parameters are summarized in Panel A of Table 3.

4.2 Simulated moments estimation

Taking as given the set of prespecified parameters described above, we then estimate the remaining structural parameters $\Theta \equiv (\delta, \gamma, \psi, \nu, \varpi, \phi_0, \phi_1, \phi_h)$ by minimizing a standard objective function:

$$\hat{\Theta} = \arg\min_{\Theta} \left(M - m(\Theta, \Theta_0)\right)' \mathbf{W} \left(M - m(\Theta, \Theta_0)\right),$$

where $m(\Theta, \Theta_0)$ is the vector of reduced-form statistics of the simulated variables, M are their empirical counterparts, and W is a weighting matrix.

For a given set of parameter values, we first solve for the optimal policies from the household problem numerically. Then, we simulate a panel of households, which are initialized by randomly drawing pairs of liquid assets a_i and mortgage balance b_i over the state space for all N households in the cross section. We use a cross section of $N = 1000$ households and compute all of the statistics m along the aggregate time path of $T = 2000$ (annual) periods, after burn-in.

Data moment targets We estimate the preference and transaction cost parameters by targeting 14 moments of the data (sources detailed in Appendix D). These include 3 unconditional means applying to the whole population: (1) aggregate ratio of nondurable and non-housing services consumption to income, (2) average household-level consumption growth volatility (based on the Consumer Expenditure Survey estimates reported by Wachter and Yogo (2010)), and (3) the average homeownership rate.

There are 6 moments relevant to the homeowner subset of the population: (4) the average ratio of liquid asset holdings to income; (5) the average ratio of household mortgage debt to income; (6) the average ratio of HELOC balances to income; (7) the average number of refinance loans relative to the number of homeowner households; (8) the average loan-to-income ratio upon refinancing; (9) dollar cash-out as a share of aggregate refinancing volume. There is also one moment for the renter population: (10) the average ratio of liquid asset holdings to income for the renter subset of the population.

All of the cross-sectional moments except (8) and (9) are based on all the 1989-2010 waves of the Survey of Consumer Finances, whereby we exclude the top 25% of households sorted on liquid assets (similarly to the approach of Gomes and Michaelides (2005)). In the data, the wealth distribution is heavily skewed to the right, which means that an "average" household in the data is not representative of a typical household that our model aims to replicate. This is in part because the restricted menu of assets that households are able to invest in as well as the preference specification that we use limit the extent of wealth dispersion that our model can generate.[16] Refinancing and cash-out loan moments (8) and (9) above are based on HMDA data.

The remaining 4 moments describe the dynamics of refinancing and cash-out behavior estimated via linear regressions of these variables on aggregate income growth and house price growth rates as described in Section 2. Table 4 reports both the target empirical moments and the simulated moments corresponding to the minimized objective function, as well as several additional moments that were not targeted in the estimation.

Since we use more moments than parameters, the model is over-identified. We use a diagonal weighing matrix that is scaled by the empirical moments in question as a normalization, that is, $\mathbf{W} = \mathrm{diag}(M)^{-1}\mathbf{S}\,\mathrm{diag}(M)^{-1}$, where $\mathrm{diag}(M)$ is a diagonal matrix with the empirical moments as the diagonal elements. The diagonal matrix $\mathbf{S}$ has elements of ones corresponding to all of the moments, except: (i) average debt balances and the refinancing rate have the weight equal 6, (ii) liquid asset holdings and average consumption growth volatility for homeowners each have the weight of 4, (iii) the 4 regression coefficients, which have the weight of 3, and (iv) the mean liquid assets of renters have the weight of 0.1. These weights reflect the fact that we are most interested in capturing the leverage and liquidity choices of homeowners. We use this pre-specified weighting matrix rather than a matrix that

[16] In our model all households are ex ante identical, and all of the heterogeneity is due to idiosyncratic shocks, which are transitory. Moreover, in our model household preferences are homothetic, while explaining the large amount of asset holdings by the wealthy households typically requires non-homotheticities, e.g. Carroll (2000), DeNardi (2004), Roussanov (2010).

Table 4: Target Moments for the Estimation and Model Outputs

Moment	Variable	Data	Model	s.e.
		Panel A. Targeted Moments		
All Households:				
1. Consumption/Income	c_i/y_i	0.66	0.71	0.01
2. Consumption growth volatility, %	$\sigma(\Delta \log c_{i,t+1})$	12.0	16.4	0.01
3. Homeownership rate, %	$E[I^h]$	66.0	67.5	0.08
Homeowners:				
4. Liquid assets/Income	a_i^+/y_i	0.30	0.24	0.04
5. Mortgage/Income	b_i/y_i	1.01	0.96	0.08
6. HELOC/Income	$-a_i^-/y_i$	0.07	0.08	0.01
7. Refinancing rate, % of homeowners	$REFI$	6.9	11.3	0.02
8. Refi loan/Income	b_i'/y_i	1.41	2.74	0.14
9. Dollar cash-out/Refi loan	$(b_i' - b_i)^+/b_i'$	0.12	0.51	0.03
Renters:				
10. Liquid assets/Income	a_i^+/y_i	0.17	0.15	0.06
Refinancing Regression:				
11. Coefficient on Z	β_Z^{REFI}	-0.25	-0.24	0.41
12. Coefficient on $\Delta \log H$	β_H^{REFI}	0.15	0.08	0.14
Cash-out Regression:				
13. Coefficient on Z	β_Z	-0.12	-0.18	0.43
14. Coefficient on $\Delta \log H$	β_H	0.06	0.11	0.15
		Panel B. Additional Moments		
Volatility of aggregate consumption growth, %	$\sigma(\Delta \log C_{t+1})$	2.7	3.9	0.01
Sensitivity of consumption to Z shocks	β_Z^C	0.46	1.30	0.20
Sensitivity of consumption to H shocks	β_H^C	0.06	0.09	0.05
Sensitivity of consumption to lagged r	β_r^C	0.07	0.13	0.43
Sensitivity of consumption to lagged R	β_R^C	0.09	0.17	0.65
Refinancing regression coefficient on R	β_R^{REFI}	-1.91	-0.96	0.67
Cash-out regression coefficient on R	β_R	-0.43	-0.59	0.73

is based on the estimated variance-covariance matrix of moments (such as the efficient GMM weighting matrix of Hansen (1982)) in order to make sure that the information in some of the economically important but relatively imprecisely estimated moments (like the regression coefficients) is not down-weighed too much.

In order to conduct statistical inference we compute the variance-covariance matrix of sample moments Ξ using simulation under the null of the model, as described in Appendix C.

4.3 Estimation results

The targeted empirical moments and their model counterparts are reported in Panel A of Table 4 along with the simulated standard errors.

In our model, the average ratio of consumption to income at 0.71 is slightly above the 0.66 in the aggregate data (using both nondurable and durable goods expenditures, as well as non-housing services); according to the model this moment is estimated very precisely, with a standard error of 1%, which implies that statistically this difference is significant, even though it is economically small. The model-implied annual household-level consumption growth volatility of 16.4% is much higher than the 9% target estimated by Wachter and Yogo (2010), which is constructed to reduce measurement error, but it is consistent with the estimate of Brav, Constantinides, and Geczy (2002) based on the CEX data (16-18% for households with total assets exceeding $2,000). The model implies an average homeownership rate of 67.4%, quite close to the 66% average homeownership rate in the data.

The 16.4% household-level consumption growth volatility is only slightly below the unconditional labor income growth volatility of 16.6%, implying limited consumption smoothing on average. The model tries to match simultaneously a low level of average liquid asset holdings, a high level of average debt holdings (both of which require low risk aversion), and a moderate consumption volatility (which requires high risk aversion). Although home equity can help homeowners smooth income shocks in bad times, the financial leverage tends to raise consumption volatility on average.

The model does a good job matching the average liquid asset holding and mortgage balances for homeowners in the data. Mortgage debt is approximately equal to annual household income on average (0.96 of income on average in the model compared to 1.01 in the Survey of Consumer Finance (SCF) data). Households pay down some of the mortgage balances over time for two reasons. First, mortgage borrowing is generally a costly way to finance consumption due to the interest rate differential between mortgage loans and personal savings. Except when the term structure of interest rates is sufficiently flat that the effective (after-tax) borrowing rate is equal to or lower than the short rate, households optimally choose to repay part of their mortgage debt rather than holding too much in liquid assets. Second, by partially repaying the mortgage debt, households can maintain some home equity "for the rainy day." Since accessing housing collateral is costly, home equity is an illiquid form of saving that can be tapped for consumption purposes infrequently, e.g., following large negative income shocks. The model also matches the average holdings of second-lien loans reasonably well (0.07 of household income in the data vs. 0.08 in the model, insignificantly different statistically given the standard error of 0.01).

Despite the low return on liquid assets, households still hold liquid assets equal to 24% of income in the model, which is close to the amount observed in the SCF data (30%). It is more efficient to use liquid assets to buffer small fluctuations in income due to the costs of accessing home equity via cash-out refinancing. Liquid assets also become highly valuable in cases when the borrowing constraints (LTV or LTI) bind.[17] The model implies a reasonable level of liquid asset holdings for renters at 15% of annual household income vs. 17% in the SCF data.

About 11.3% of homeowners per year refinance their mortgages in the model, compared to 6.9% in the SCF data. The average loan-to-income ratio for the new loans originated

[17] Using 2004 SCF data, Vissing-Jørgensen (2007) estimates that by using their lower-return liquid assets to accelerate the repayment of higher-cost housing debt U.S. consumers would have saved $16.3 billion - see discussion in Guiso and Sodini (2013). Telyukova (2013) analyzes the role of liquidity in explaining the related puzzle of concurrent credit card debt and savings account holdings documented by Gross and Souleles (2002), while Laibson, Repetto, and Tobacman (2003) argue that consumer self-control problems may be necessary to explain quantitatively the extent of the puzzle.

from refinancing in the model (2.74) is significantly higher than the average value in the 2001 SCF (1.41) and the HMDA data for 1993-2009 (1.90). Accordingly, the amount cashed out conditional on refinancing is also high, equaling to 51% of new loan balances, compared to 12% in the data. Estimates from the data are based on the average cash-out share of refinance originations for prime, conventional loans, and average loan-to-income (for all refinance loans). To the extent that these estimates are representative of the U.S. homeowners, the model predicts too much cash-out as well as too frequent refinancing into large mortgages in general, with the differences being both economically and statistically significant. It is a challenge for the model to simultaneously match the refinancing rate and the dollar amounts of cashed-out home equity. While raising the fixed cost of loan origination helps reduce the frequency of refinancing, it makes households cash out even more each time they refinance.

On the set of moments from the refi and cash-out regressions, the model matches the signs and approximately the magnitudes of all the coefficients on income growth (β_Z) and on house price growth (β_H), especially in the case of cash-out regression. Both the refinancing rate and the dollar cash-out to income ratio comove positively with house price growth, and negatively with income growth, as we find in the data. While these regression coefficients are estimated quite imprecisely, as evidenced by the large standard errors that we report, targeting these coefficients is important for capturing the cyclical dynamics of household demand for liquidity, which helps to identify some of our key structural parameters.

Next, the estimated values of the preference and transaction cost parameters are reported in panel B of Table 3, accompanied with the standard errors in the parentheses. The preference parameters implied by the moments above are the subjective discount factor $\delta = 0.920$, the coefficient of relative risk aversion $\gamma = 3.036$, and the intertemporal elasticity of substitution $\psi = 0.301$. These parameters imply a moderate degree of risk aversion and a limited willingness to substitute consumption intertemporally, i.e. a desire for a smooth consumption profile over time. These parameter estimates are driven largely by the low target level of liquid asset holdings, high debt levels, and the observed sensitivity to changes in interest rates and economic conditions embedded in the refinancing frequency and the regression coefficients. In particular, our estimate of the IES is close to the estimate obtained by Vissing-Jørgensen (2002) using stockholder household-consumption data from the CEX (0.299).[18]

While a number of studies that estimate the IES using the aggregate log-linearized Euler equation following Hall (1988) find values very close to zero, such an approach would not be valid in an economy that conforms to our model, given the substantial heterogeneity and frictions.[19] As Table 4 Panel B reports, the estimated slope coefficient from the regression of consumption growth on the lagged risk-free rate based on the simulated data from the baseline model is only 0.13, while the coefficient from the regression of consumption growth on the lagged long-term mortgage rate R is 0.17, both close to one half of the true value of the IES. This low sensitivity of consumption growth to lagged interest rates is largely due to the presence of long-term mortgages, which are costly to to adjust, as we demonstrate in Table 5 below. In addition, the very presence of frictions (i.e., the borrowing constraints) makes consumption much less sensitive to the fluctuations in interest rates, potentially muting the impact of monetary policy.

The estimated implied average rent/income ratio parameter is $\varpi = 1.324$. This parameter is identified jointly by the average consumption-income ratio and the share of homeowners as well as the balance sheet moments, since the benefit of homeownership is in large part the avoidance of rental expenses but also the asset and collateral value of housing.

Households use debt primarily as a way of smoothing consumption and financing new home purchases. Existing debt balances are refinanced either to reduce the coupon rate k,

[18] Our estimate of the IES differs from values typically used to reconcile asset pricing facts with consumption dynamics in representative-agent models. For example, Bansal, Kiku, and Yaron (2012) estimate IES of around 2 using aggregate consumption and asset price data, while their estimate of the coefficient of relative risk aversion is twice as large as ours. This is not surprising since the only risky asset that we target in the data is housing (and mortgage). Moreover, we target households in the bottom 80% of the wealth distribution, who exhibit low rates of stock market participation. Vissing-Jørgensen (2002) obtains estimates of the IES above one for households in the upper tail of the wealth distribution who participate in financial markets; see also Attanasio and Weber (1995) and Vissing-Jørgensen and Attanasio (2003).

[19] Carroll (2001) and Hansen, Heaton, Lee, and Roussanov (2007) discuss some of the issues associated with the standard approaches to estimating the IES.

or to cash-out equity. The quasi-fixed and proportional costs of refinancing, ϕ_0 and ϕ_1, are primarily identified by targeting empirically observed average refinancing rates, in terms of both frequency and loan size. They are also influenced by the average level of mortgage debt, since higher transaction costs make higher balances less attractive by effectively lowering the value of the refinancing option, as well as by making home-equity withdrawal via cash-out more expensive. Anecdotal evidence suggests that explicit costs of roughly $2\% - 5\%$ of loan amount are paid when refinancing a mortgage loan of average size, in addition to non-pecuniary information processing costs and the opportunity cost of time required to process the transaction. In the estimation, we obtain a quasi-fixed cost of 15.4% of permanent income (or 3.9% of the house value on average) and a proportional cost of 1.4%, which is comparable to the costs calibrated by Campbell and Cocco (2003).[20]

The model implies that the cost of buying (or selling) a house ϕ_h is 13.5% of the house value. This parameter is identified primarily by the average homeownership rate but also by the asset holding levels among homeowners and renters, since this parameter controls the cost of transition from one group to another. This estimated cost is high, although it is meant to capture the psychic and physical costs of moving, besides the actual pecuniary transaction costs (such as transfer taxes and realtor commissions).

As indicated by the standard errors, most of the parameters are estimated fairly precisely in the sense that the sampling uncertainty about the data moments, under the null of the model, translates into tight confidence bands for the point estimates. All of the parameters are statistically significantly different from zero. The discount factor δ is statistically significantly lower than unity. Interestingly, the coefficient of relative risk aversion cannot be distinguished from the inverse of the IES, suggesting that the standard separable utility function with constant relative risk aversion provides a reasonable description of household preferences.

Finally, Panel B of Table 4 reports several moments that are not targeted in the structural

[20] Empirically the bulk of explicit cost of refinancing can be attributed to title insurance, which is proportional to house value, whereas the non-monetary costs such as the opportunity cost of time spend searching for an attractive mortgage rate and preparing the necessary documents are likely quasi-fixed.

estimation. Checking the ability of the model to match these moments is a form of out-of-sample test. The volatility of aggregate consumption growth in the model is 3.9%, compared to 2.7% in the data. This higher volatility of consumption is driven largely by its greater sensitivity to fluctuations in aggregate income than observed in the data, with the regression coefficient on Z being 1.3 vs. its counterpart in the data of 0.46. As mentioned above, aggregate consumption is not very sensitive to fluctuations in interest rates. At the same time, it is quite responsive to changes in house prices, with the regression coefficient of aggregate consumption growth on the house price growth of 0.09 (albeit with a relatively wide standard error of 0.05), compared to the estimated coefficient of 0.06 in the U.S. data. While the latter is difficult to estimate precisely in the aggregate data, using disaggregated data over the 2006-2009 period, Mian, Rao, and Sufi (2013) estimate average marginal propensities to consume out of housing wealth between 0.05 and 0.11. While we only target the cyclical behavior of refinancing and cash-out in our estimation, the model matches reasonably well the sensitivities of the total refinancing rate and dollar cash-out to the fluctuations in the mortgage rate. In the refinancing regression, the coefficient on mortgage rate, β_R^{REFI}, is -0.96 in the model, compared to -1.91 in the data. In the cash-out regression, $\beta_R = -0.83$ in the model vs. -0.43 in the data.

4.4 Effects of Labor Income Risk and Financing Constraints

In order to further examine the model's mechanism, we present a range of comparative statics. In particular, here we focus on the effects that counter-cyclical labor income risk and financing constraints have on household consumption and financing decisions. We report the simulated moments from the model for each of the model specification alongside the baseline that uses the estimated parameter values.

Labor income risk In Table 5, column (1) reports the baseline model results. In column (2), we shut down heteroscedasticity in the idiosyncratic labor income process by setting $\sigma(Z_G) = \sigma(Z_B) = 8\%$. The removal of the counter-cyclical variation in labor income risk (and lower average level of labor income risk) significantly reduces the benefit of homeownership through home equity savings. As a result, homeownership drops from 67.4% to 46.3%. Homeowner households become more aggressive in taking on leverage (mortgage-to-income ratio rises from 0.94 to 1.63). High leverage implies lower levels of home equity savings for homeowners. Moreover, both homeowners and renters hold less in liquid assets. In addition, homeowners refinance more frequently (refinancing rate doubles from 11.1% to 22.6%) and respond more strongly to interest rate fluctuations (with β_R^{REFI} changing from -0.96 to -1.98), while cash-out becomes much less sensitive to changes in aggregate income (with β_Z changing from -0.18 to -0.01). Note that the baseline model produces virtually no defaults, while the homoscedastic model features about 0.1% default rate, on average. This is because the homoscedastic model features less "tail risk" in labor income shocks and, as a consequence, less conservative leverage choices than the benchmark model, which lead to more frequent defaults. While this frequency of default may appear low relative to the data, Gerardi, Herkenhoff, Ohanian, and Willen (2018) show that in fact the vast majority of households with negative equity in their home do not default, even after experiencing a large negative income shock.

Collateral vs. debt service constraints Specifications (3) and (4) consider the effects of relaxing the collateral constraint (LTV) and debt service constraint (LTI). We relax the LTV constraint by setting $\xi_{LTV} = 100\%$ and remove the LTI constraint by setting $\xi_{LTI} = \infty$, respectively. These comparative statics capture the perception that mortgage lending standards were dramatically relaxed over the course of the housing boom. The relaxation of the LTV constraint can also mimick the Homeowner Affordable Refinance Program (HARP) instituted by the U.S. government in 2011, which was intended to allow certain underwater homeowners to refinance.

Relaxing the collateral constraint leads to a simultaneous increase in leverage (in terms of mortgage-to-income ratio) and liquid asset savings for homeowners. The net effect is that consumption volatility becomes slightly higher. Homeownership increases from 67.4% to

Table 5: Effects of Labor Income Risk and Financing Constraints

	(1) Baseline	(2) $\sigma = 8\%$	(3) $\xi_{LTV} = 1$	(4) $\xi_{LTI} = \infty$	(5) (3) + (4)	(6) $\underline{a} = 0$	(7) $b = 0$
All Households:							
Cons. growth vol, %	16.6	21.1	16.9	16.9	17.0	16.6	21.1
Homeownership rate, %	67.4	46.3	74.7	81.9	88.4	61.6	75.4
Homeowners:							
Liquid assets/Income	0.24	0.16	0.27	0.27	0.29	0.34	0.05
Mortgage/Income	0.94	1.63	1.10	1.56	2.07	0.91	-
HELOC/Income	0.09	0.08	0.08	0.08	0.08	-	1.06
Refinancing rate, %	11.1	22.6	11.3	15.8	17.6	11.8	-
Refi loan/Income	2.73	2.26	2.87	3.19	3.65	2.71	-
Cash-out $/Refi loan	0.51	0.34	0.50	0.40	0.34	0.45	-
Default rate, %	0.0	0.1	0.4	0.6	2.1	0.0	0.3
Renters:							
Liquid assets/Income	0.15	0.07	0.19	0.09	0.07	0.15	0.17
Refinancing Regression:							
Coefficient on R, β_R^{REFI}	-0.96	-1.98	-1.10	-1.30	-1.56	-0.83	-
Cash-out Regression:							
Coefficient on R, β_R	-0.59	-0.17	-1.10	-0.22	-0.18	-0.45	-
Coefficient on Z, β_Z	-0.18	-0.01	-0.31	-0.45	-0.52	-0.20	-
Coefficient on H, β_H	0.11	0.19	-0.05	0.37	0.23	0.10	-
Aggregate Consumption:							
Growth volatility, %	3.9	4.0	4.0	4.6	5.0	3.7	5.0
Sensitivity to Z, β_Z^C	1.30	1.20	1.35	1.36	1.50	1.24	1.58
Sensitivity to H, β_H^C	0.09	0.16	0.09	0.17	0.16	0.10	0.04
Sensitivity to lagged r, β_r^C	0.13	0.08	0.14	0.13	0.16	0.10	0.28
Sensitivity to lagged R, β_R^C	0.17	0.11	0.19	0.16	0.19	0.13	0.22

Note: Refinancing rate and default rate are in percentage of homeowners.

74.7%. This is because (1) more "marginal" households are able to enter the housing market with easier access to credit, and (2) the benefit of homeownership is higher when households can access their home equity savings to a fuller extent. A loosened LTV constraint also raises the sensitivity of cash-out to aggregate mortgage rate and income shocks. Interestingly, the relation between cash-out and house prices changes significantly from the baseline case (β_H changes sign from 0.11 to -0.05). Thus, households cash-out more, not less, following drops in house prices. Two effects are at work in determining how cash-out responds to house price shocks. First, a rise in house price relaxes the LTV constraint, which helps generating a positive relation between cash-out and house price changes. Second, to the extent that house price drops are associated with a deterioration of the state of the economy, the demand for extracting liquidity from home equity becomes stronger. This effect generates a negative relation between cash-out and house price changes. If the LTV limit is relatively low, as in the baseline case, the former effect dominates. If the LTV limit is high, meaning the collateral constraint is already relatively slack, as in specification (3), the latter effect dominates.

Next, in specification (4) we remove the LTI constraint ($\xi_{LTI} = \infty$). Similar to the case where we relax the LTV constraint, removing the LTI constraint raises homeownership, average mortgage balances, and liquid asset holdings for homeowners. As a result of large mortgage balances, refinancing becomes more frequent and more sensitive to interest rate changes. Moreover, households cash out significantly more than in the baseline case following aggregate income shocks (β_Z changes from -0.18 to -0.45).

Besides these similarities, there are important qualitative differences in the effects of relaxing the LTI vs. LTV constraint. Relaxing the LTI constraint is particularly relevant for low-income households. These households can now become homeowners without significant savings, which explains why renters on average hold less liquid assets in specification (4) than in the baseline case (opposite to specification (3)). The removal of the LTI constraint also allows more low income households to access home equity savings. With these households accounting for a larger part of aggregate cash-out activities, aggregate cash-outs become less sensitive to mortgage rates and more sensitive to aggregate income shocks. Cash-outs now respond more strongly to house price changes (β_H rising from 0.11 to 0.37), which is opposite to what happens with the relaxation of the LTV constraint. This is intuitive, since a relaxation of the LTI constraint means more households can cash out after a rise in house price relaxes the LTV constraint.

In specification (5), we examined the combined effect of setting maximum LTV to 100% and removing the LTI limit. This change has a dramatic effect on almost all of the moments, illustrating how the two constraints reinforce each other. The amount of risk in the economy increases despite the greater ability to smooth fluctuations, which is due to the endogenous response of households of choosing greater leverage and higher investment in (risky) housing. While household-level consumption volatility increases only slightly, to 17% per annum, aggregate consumption growth volatility is the highest among all specifications, at 5%, with sensitivities to all of the aggregate state variables displaying increases. Most notably, the rate of mortgage default is also the highest, at 2.1% of homeowners per annum. With higher leverage it is more likely that a household would find its home equity negative after a decline in house prices, which is a necessary (but not sufficient) condition for a strategic default to be optimal.[21] The results here suggest that the simultaneous relaxation of the LTV and LTI constraints can have a dramatic (more than additive) effect on default.

Long-term mortgage vs. HELOC Finally, we compare the different effects that long-term mortgages and HELOCs have on households. Such a comparison is important since short-term debt is featured in most macroeconomic models of the housing collateral for tractability, while long-term mortgages account for the majority of household debt in the data.

In specification (6), we remove HELOCs completely by setting $\underline{a} = 0$. Since accessing home equity through HELOCs does not incur any refinancing costs, removing HELOCs effectively tightens the borrowing constraints and reduces the attractiveness of homes as a saving vehicle. A direct effect is reduced homeownership, from 67.4% in the baseline case to 61.6%. Another consequence is that homeowner households now hold more liquid assets (the liquid assets to income ratio rises from 0.24 to 0.34) and less debt (the total debt to income ratio falls from 1.03 to 0.91). As discussed before, HELOCs are used mainly to smooth small idiosyncratic income shocks. Without HELOCs as a liquid source of credit, households simply substitute into liquid assets, while their consumption and mortgage financing behaviors are

[21] Corbae and Quintin (2013) analyze the effect of the loosening and subsequent tightening of leverage constraints on mortgage default following the decline in house prices; see also Campbell and Cocco (2015) for a detailed analysis of household default decisions in the presence of labor income shocks and different mortgage products.

not significantly affected.

In specification (7), we remove long-term mortgage contracts and consider a case where households can borrow via unrestricted HELOCs, limited only by the LTI and LTV constraints (HELOCs are capped at 30% of permanent income in the baseline case). A key difference from long-term mortgages is that HELOCs are subject to the LTV and LTI constraints each period, as opposed to only when households refinance or obtain a new mortgage. For example, following a drop in housing prices that reduces its home equity below 20% of house value, a household must pay down enough of its short-term debt to satisfy the LTV constraint upon roll-over or be forced to default. Similarly, a drop in income may make the LTI constraint binding, forcing household to cut consumption to avoid default. Consequently, this risk of forced deleveraging makes short-term mortgages less useful as tools of consumption smoothing than long-term mortgages.

Indeed, the simulated data for this case features higher volatility of consumption growth, at both the individual and the aggregate level (at 21% and 5%, respectively) and greater sensitivity of consumption to aggregate income shocks (at 1.58 vs. 1.3 in the baseline case). This is despite of the fact that the total debt level is comparable to the baseline case (total debt to income ratio is 1.06 vs. 1.03 in the baseline case). There is also a rise in default rate to 0.3%. However, our assumptions on transactions costs (i.e., costless access to HELOCs) make these short-term loans attractive, especially for new homeowners (or movers) who effectively face lower moving costs than under our baseline case. Consequently, the homeownership rate is higher, at 75%. In all, our results show that while models of short-term and long-term debt may be hard to distinguish on the basis of aggregate leverage alone, they have rather different implications for the ability of households to insure against aggregate and idiosyncratic shocks, and for the composition of household assets in terms of their relative liquidity.

5 Cross-Sectional Evaluation

Having examined the aggregate moments of the estimated model, we now turn to its implications for the dynamics of household financing and consumption in the cross-section of households. We focus on the behavior of homeowners with respect to their use of mortgage debt as a key tool of household risk management.

5.1 Leverage and Refinancing in the Cross-Section

The main focus of our paper is households' use of mortgage refinancing as a liquidity management tool. Thus, we begin by evaluating our model's ability to explain refinancing behavior in the cross-section of households over time. A key dimension of household heterogeneity in our model is household income and in particular its transitory component. In Figure 3, we confront the model's cross-sectional predictions with the empirical evidence, which are based on data from SCF for years 1998 through 2010. In the model as well as in the data, we sort households into quintiles based on income divided by house value ($\tilde{y}_{it}/h_{i,t}P_t^H$) and on the leverage ratio (measured by the ratio of mortgage debt to income); in both cases we sort households in the panel (i.e. pooling the time-series and the cross-section), which allows the model to capture the changing composition of refinancing households over time.

The model matches the cross-sectional distribution of mortgage debt-to-income ratios ($b_{i,t}/y_{i,t}$) remarkably well (Panels A and B). The debt-to-income ratio declines monotonically with income both in the model and in the data. In the model, the bottom income quintile on average has mortgage balances of about twice the annual income, while the top income quintile has average mortgage balances of about a quarter of the annual income. In the data, the mortgage debt-to-income ratios also decline with income, but are somewhat higher in value within each quintile. This is in large part due to the fact that the model imposes a hard LTI constraint, which may not have been imposed in reality on some of the households (especially during the housing boom period).

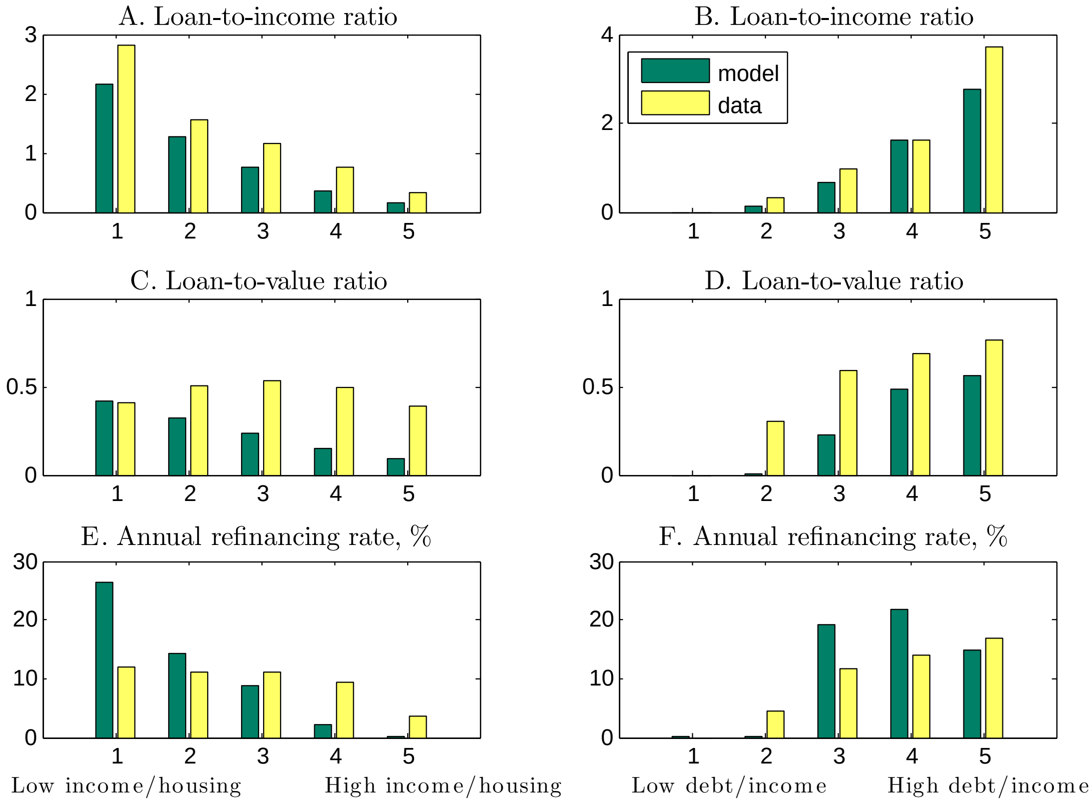

Figure 3: **Cross-sectional implication: Model vs. Data.** Model estimates are based on a pooled simulated panel; data analogs are based on pooled SCF data for years 1998-2010.

The model also does a reasonable job capturing the empirical distribution of loan-to-value ratios (LTV) across households sorted on leverage (Panel D). Those households in the bottom 40% of the LTV distribution have essentially zero debt in the model and in the data, and both increase monotonically to about 0.5 in the model and 0.7 in the data. The model has a harder time matching the empirical distribution of LTV across households sorted on income (Panel C). In the data, LTV is slightly hump-shaped in income-to-house ratio. In the model, the average LTV decreases in income from 0.4 to about 0.1 as normalized income increases. The hump-shaped pattern of LTV in the data is likely due to the mechanical effect that higher income-to-house ratio tends to be associated with lower house value and thus higher LTV. This effect is weaker in the model because house prices (per unit of housing) are common for all households at any point in time, whereas in the data there is significantly more variation in house prices at regional and even individual level.[22]

Finally, the model matches the distribution of refinancing rates fairly well on both dimensions. In the model, refinancing is concentrated among households whose incomes are low relative to the value of their houses (as shown in Panel E), which allows them to borrow against their home equity to smooth consumption over time. A similar (roughly monotonic) decreasing pattern is evident in the data, albeit the relationship is much less steep. While some of this variation is driven by idiosyncratic income shocks, transitory variation in house prices also contributes by determining the amount of home equity that can be cashed out.[23] The higher the level of mortgage debt, the greater is the incentive to refinance (into a lower rate); at the same time, highly levered households are more likely to face binding borrowing constraints that prevent them from extracting home equity. It is evident that both effects are at work in the model, although medium leverage households in the model tend to refinance more frequently than in the data (see Panel F).

5.2 Drivers of Refinancing

The results about the leverage and refinancing patterns in the joint distribution of household income and mortgage debt in Figure 3 are consistent with our model's prediction that liquidity demands are a key driver of refinancing. In order to further examine the two drivers of refinancing – reducing financing costs and home equity extraction, we need to be able to observe household balance sheets over time, i.e. before and after refinancing. We utilize Panel Study of Income Dynamics (PSID), which allows us to track households over time, before and after they refinance their mortgages. We pool all waves between 1999 and 2013

[22]In the Appendix we present additional empirical evidence that both the amount of refinancing activity and the average LTI on refinanced loans vary negatively with income and positively with house price appreciation at the state level.

[23]It is well known that when interest rates fall, wealthier/higher income households are more likely to take advantage of the opportunity to refinance (e.g., Fuster and Willen (2011)). Even though we estimate fairly large quasi-fixed costs of mortgage origination, our model generates too much refinancing for low income households relative to the data, compared to the high-income households, which could at least in part be due to cognitive costs associated with understanding the refinancing process are decreasing with household income. See e.g., Woodward and Hall (2010).

and identify refinancing if mortgage origination dates between neighboring waves differ and the most recent origination year does not precede that of the earlier date, while household owns the same home in the two waves.

Figure 4 presents information on the refinancing behavior across households sorted on normalized income and leverage (again measured by mortgage debt to income). While the model overstates the average amount of home equity extraction via cash-out refinancing (as already shown in Section 4.3), it is still worthwhile to examine the cross-sectional variation in refinancing incentives when the cash-out mechanism is operative. Thus, in Panel A we plot the distribution of cash-out to income ratio, normalized by its average value. In the model, liquidity needs drive much of the refinancing behavior. Consequently, conditional on refinancing, the average dollar cash-out to income ratio is decreasing in income, from close to 1.5 in the bottom income quintile to about 0.25 in the top. This monotonic pattern is also observed in the data, although there is less variation across the top income quintiles than predicted by the model. As Panel B shows, the cash-out to income ratio is also decreasing in household leverage in the model, which is largely due to the LTI constraint. It is highest in the bottom two leverage quintiles (households who go from essentially zero debt all the way up to the constraint) and falls monotonically across the top three quintiles. Here the model fails to match the data, which features a somewhat U-shaped, rather than decreasing, pattern of cash-out as a function of leverage.

Panels C and D plot the the ratio of the new mortgage rate upon refinancing k' to the old rate k. In the model, this ratio is above unity on average for the bottom three income quintiles but significantly below unity for the top income quintile. This result is a clear indication that liquidity demands drive much of the refinancing activity for low income homeowners. These households are willing to increase their average debt service cost in order to access liquidity. In contrast, high income households tend to have lower mortgage balances. They will require a significant drop in mortgage rate to be willing to incur the fixed cost for refinancing. In contrast, the rate ratio declines with the debt to income ratio across the top three quintiles.

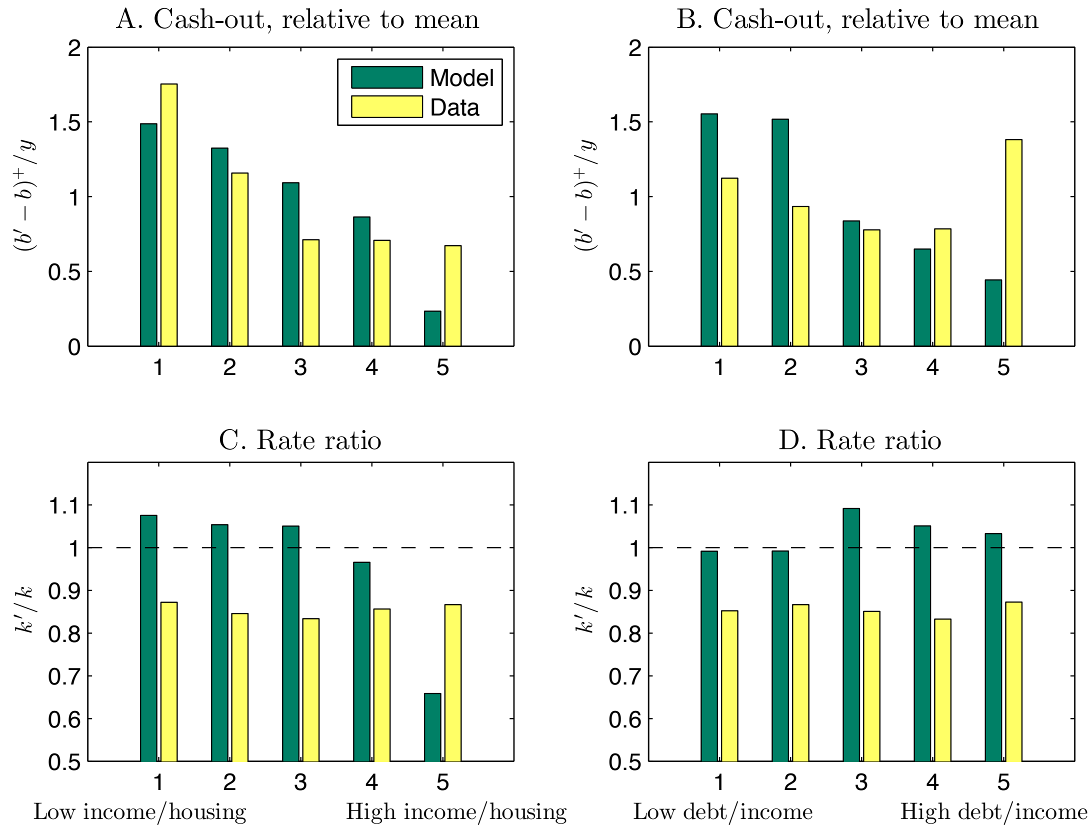

Figure 4: **Refinancing for the cross section of households.** Model estimates are based on a pooled simulated panel; data analogs are based on pooled PSID data for years 1999-2013.

This is because (1) larger debt balances make refinancing into a higher rate more costly; (2) due to the LTI constraint, the amount of liquidity that households can access through refinancing drops as leverage rises. Interestingly, in the data the ratio is instead essentially flat across both the income and leverage quintiles, and is on average below unity. However, there is substantial variation in the rate ratio over time, as is evident in Panel B of Figure 1. Furthermore, this ratio is well above one and somewhat U-shaped across both the income and the debt-income quintiles during the period that saw a boom in cash-out refinancing. Overall, our model seems to understate the interest-savings motive for mortgage refinancing (relative to the liquidity motive) for the lower-income and high-leverage households, in part due to the fact that we target the cyclical behavior of refinancing and cash-out activity, which is our

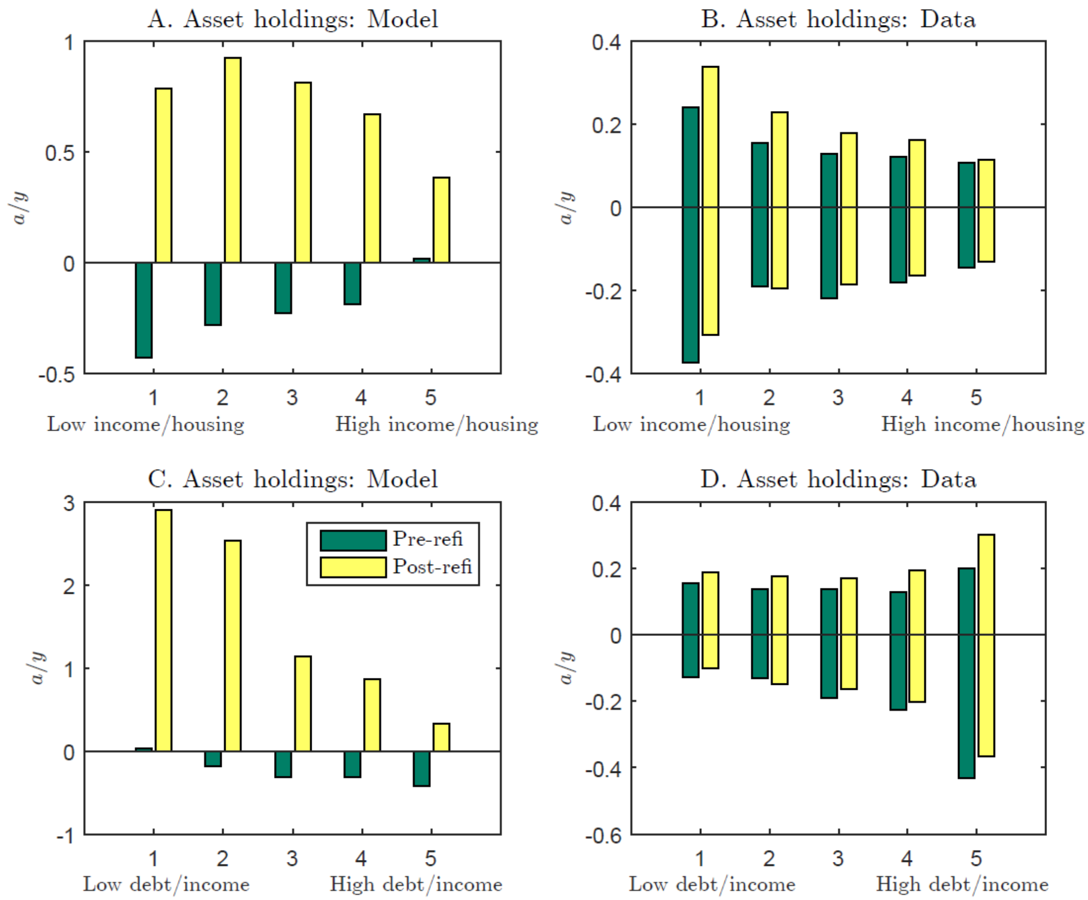

Figure 5: **Refinancing and liquid assets: before and after.** Model estimates are based on a pooled simulated panel; data analogs are based on pooled PSID data for years 1999-2013.

main focus, rather than the sensitivity of refinancing to interest-rate movements, which is the standard channel considered in the literature.

Figure 5 plots households' average HELOC ($a < 0$) and liquid asset holdings ($a > 0$) before and after refinancing, both in the model and in the data. As Panel A shows, in the model, refinancing households in the first four income quintiles on average have nonzero HELOC balances before refinancing, while those in the top income quintile have a small amount of liquid assets. This suggests that liquidity-constrained households first borrow using short-term HELOCs, which have no transaction costs, before switching to cashing out home equity when the liquidity needs become sufficiently strong. After refinancing, the cashed-out home equity not only helps pay down the HELOC balances, but substantially boosts the liquid asset positions, which ranges from 80% of annual income for the bottom quintile to 40% for the top quintile.

Panel B displays the corresponding magnitudes of (safe) liquid assets (positive) and short term liabilities (negative assets) in the PSID data for the households refinancing their mortgage during the sample period. Unlike in the model, households can simultaneously hold short term debt and liquid assets in the data. While the levels of (both negative and positive) assets are smaller in the data, the general pattern matches that predicted by the model: households in the low income quintiles enter refinancing with negative liquid assets (i.e. junior and unsecured debt exceeding liquid savings), and exit with positive liquid assets (about 10% of annual income on average). This is consistent with the model's prediction that households experiencing adverse income shocks first use debt sources that are less costly to access, but upon refinancing a mortgage uses extracted home equity to extract more than enough liquidity to repay such debt.

While in the model high leverage households also use the cashed-out funds to repay HELOC balances, Panel C shows that the LTI constraint severely limits the amount of liquid assets they can raise through refinancing. For example, the liquid assets for the top leverage quintile is just above 30% of annual income after refinancing, in contrast to close to 3 times annual income for households in the bottom leverage quintile. Again, as Panel D shows, the magnitude of the changes in liquid assets as a result of cash out is much smaller in the data, especially for low-leverage households. However, the data are consistent with a key aspect of the model's prediction: highly levered households enter refinancing with substantial net debt in addition to their first mortgage, on the order of one fifth of annual income (both in the model and in the data), and exit with nontrivial liquid assets in excess of such debt (or having repaid it).

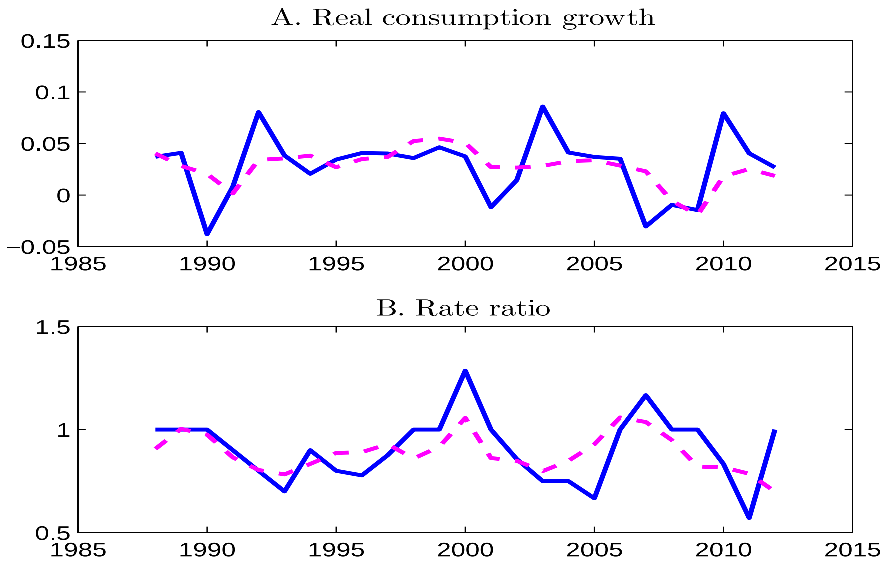

Figure 6: **Model-implied aggregate time series.** This figure plots the model-implied aggregate time series (solid lines) of real consumption growth and the median rate ratio of refinance loans and their data counterparts (dashed lines).

5.3 Historical time series

To evaluate the model's ability to match the observed history of household consumption behavior, we simulate a panel of 1000 households, who face idiosyncratic labor income shocks as well as the time series of realized shocks to the exogenous state variables in the data for the period 1988–2012.

Figure 6 Panel A depicts the annual series of real consumption growth from the model and the data. The model overstates the fluctuations in consumption growth in 1990-1991 (both the recession-induced drop and the subsequent recovery); it matches closely the rapid and smooth growth in consumption boom in the late 1990s and somewhat exaggerates the "consumption boom" of mid-2000s; it captures the large consumption drop during the Great recession, and somewhat overshoots the subsequent recovery.

Even with the empirical processes for aggregate income and house prices that we feed into the model, households inside the model can still endogenously adjust their decisions on consumption, savings, homeownership, and mortgage refinancing. The fact that our model is able to match the key consumption patterns in the data indicates that the model has done a decent job overall in modeling the endogenous household decisions. Specifically, the model captures the relaxation of liquidity constraints due to the rise in house prices in the 2000s, which allowed households to rationally withdraw home equity via cash-out refinancing (and second-lien borrowing), driving up household leverage and generating (in part) the consumption boom of the mid-2000s. The fall in house prices and income starting in 2007 following the dramatic expansion of leverage tightened households' balance sheets, causing a sharp and protracted consumption drop.

In particular, our model is able to capture the liquidity-driven refinancing activities in the data. This feature is apparent from Panel B of Figure 6, which depicts the median ratio of the mortgage rate obtained as a result of refinancing to the rate on the original (prepaid) loan. The model matches the peaks when the rate ratio goes above unity, capturing the effect of liquidity demand by constrained households at the onset of the two most recent recessions. Moreover, the rate ratio series appear to be moving in the opposite direction of the consumption growth plotted in Panel A, suggesting that absent the opportunity to refinance (and cash-out) consumption would fall even more in recessions.

5.4 The Housing Boom and Bust

The aggregate patterns in the historical time series mask substantial heterogeneity in households' responses to the aggregate shocks. What are the differences in household behavior during the housing boom from 2001 to 2006, as well as following the housing bust of 2007 and during the ensuing Great Recession? We use the simulated artificial panel based on the aggregate historical time series (as described in Section 5.3) to analyze the model's cross-sectional implications in this period. Figure 7 plots several key variables aggregated over groups of households in the model: the top (dashed line), middle (solid line), and bottom (dash-dotted line) quintile based on the debt-to-income ratio in 2001. We plot the simulated series for the years 2001-2012 to illustrate the heterogeneity in households' responses to

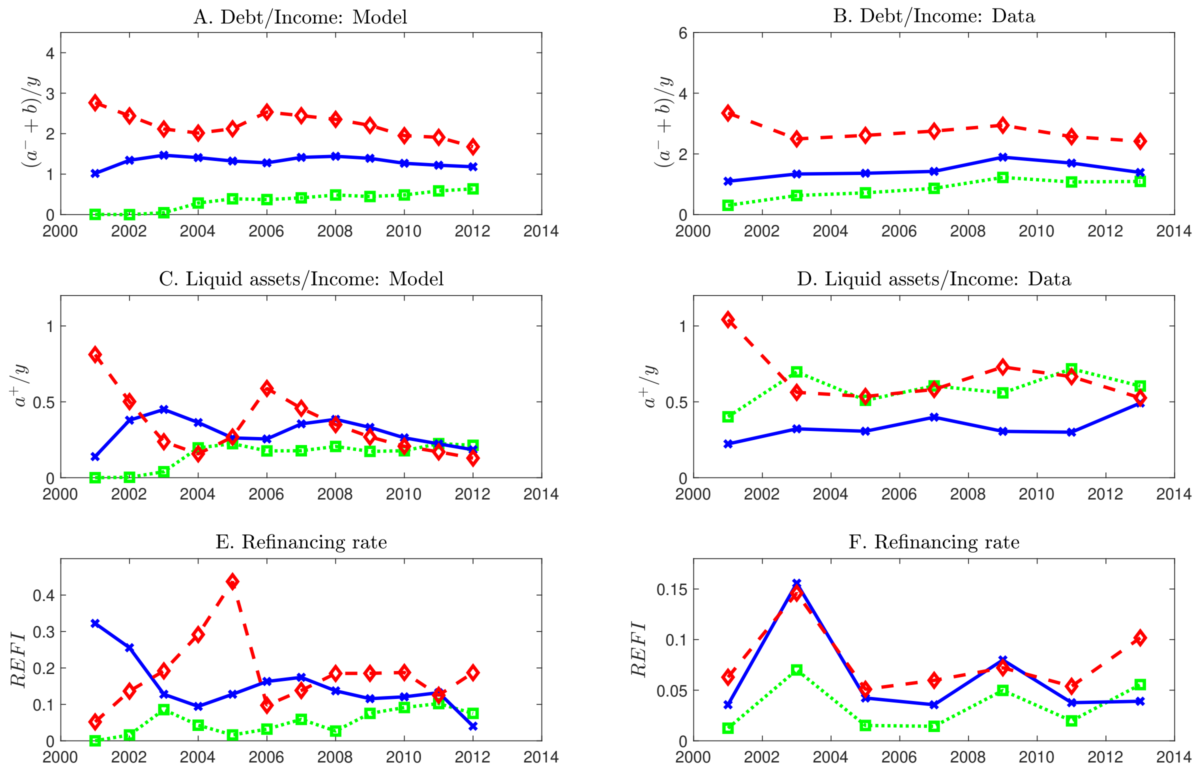

Figure 7: **Household balance sheets and refinancing, 2001-2013: model vs. data.** The dash-diamond, solid-cross, and dot-square lines represent, respectively, the top, middle, and bottom quintiles of the distribution of debt-to-income ratio in 2001. The left panels are based on model simulation, while the right panels are based on PSID data.

aggregate economic conditions. We compare it with the empirical evidence using PSID data by sorting sample households into quintiles based on their ratio of mortgage debt to annual income in 2001 for all of the waves (2001, 2003, 2005, 2007, 2009, 2011, and 2013) in which they are present in the sample.

The top two rows of Figure 7 plot the evolution of liquid assets and debt for the three groups as predicted by our model and observed in the PSID data (panels A and B, respectively). A key implication of the model is that household leverage (debt-to-income ratio) is highly persistent, which is also a salient feature in the data (see Panels C and D). Households in the top quintile of debt-to-income ratio in 2001 continue to have higher (and more volatile) leverage ratios than the other groups throughout this sample period. After an initial drop in the leverage ratio from 2001 to 2004, this group of households are the most aggressive in the model to refinance their mortgages following house price increases (as shown in panel E), which results in their leverage rising back to high levels. The leverage ratio of the top quintile also climbed back up in the data, although the climb was slower and lasted longer in the data than in the model. The patterns of refinancing are broadly matched in the data, with two important exceptions (to the extent these can be reliably detected given a relatively small number of refinancers in each wave of the PSID). First, the model overstates the refinancing of the middle quintile of borrowers in the year 2001 by a large amount. This is likely due to the fact that our discretization of the state space implies that the mortgage rate drops more dramatically that year in the model than in the data (see Figure 2). Second, the model generates a delayed refinancing wave for the most levered households, with significantly higher refinancing rates in 2004 and, especially, in 2005 (as opposed to peaking in 2003 in the data). These households were prevented from refinancing in the earlier period by the borrowing limits (including the 80% LTV limit on all borrowers) in the model. They caught up on the refinancing wave later after rising house prices helped relaxing the LTV constraints.

Not surprisingly, the model predicts that the high-leverage group will start with large liquid asset holdings (proceeds from recent cash-outs) relative to their incomes, which is also true in the data (see Panels A and B). During the 2001-2003 recession and recovery period the most levered households (typically, those who had suffered from large negative idiosyncratic income shocks during the recession) rely on these assets to support consumption, as their income slowly recovers. The ensuing house price increase presents them with an opportunity to extract additional equity out of their homes, driving a wave of refinancing in 2004-2005. We do not observe such a wave in the PSID data, where refinancing activity peaks in 2003 for all groups of households, while our model matches this peak only for the less-levered groups, likely because the borrowing constraints that we impose in the model are tighter than those prevailing during that period in the data. However, our model matches the levels of indebtedness for all groups at the onset of the Great Recession rather closely. As a result of cashing out over the boom period, the high-leverage households in the model hold significant amount of liquid assets in 2006-2007. In contrast, the other groups' liquid assets are only about a third of the high leverage group (in the data, the least-levered quintile also has a high level of asset holdings throughout much of the period, which is driven by equity and other risky financial assets that are absent from our model).

This endogenous correlation between leverage and liquid asset holdings is important for assessing the impact of income shocks on consumption. Ignoring such links can lead one to overstate the "deleveraging effect" for aggregate consumption and consumption of high-leverage households. As our model shows, during the recession, high- and medium-leverage households draw down their liquid assets over time, while low-leverage homeowners accumulate liquid assets due to elevated income uncertainty. The high-leverage households also significantly reduce their leverage over 2007-2010 as a result of debt repayment and (later) the rebound in income. In the model, they start out with mortgage debt that is 3 times the size of annual income, falling to 2 times income by 2012 (Panel C). In the data, the high leverage decile enters with debt levels close to 4 times income, reducing the debt to 3 times income by 2013 (Panel D). One of the reasons mortgage debt is higher in the data than in the model in the extremes of the distribution, especially for the most levered households, is that our model misses an important source of cross-sectional variation in the data, namely regional differences in house price appreciation.

Recent empirical work emphasizes the role of high leverage at the onset of the Great Recession on depressing household consumption, especially by the most indebted household in the areas that experienced where house prices had appreciated the most in the preceding years. Table 6 displays the cumulative consumption growth during the Great Recession for the most highly levered quintile of households and for the median household in the model and compares it to the empirical counterparts from the PSID based on Dynan (2012), who reports non-housing consumption expenditure growth from 2007 to 2009 (given our discretization of the aggregate productivity series, households begin to experience a large drop in income one year earlier in the model than in the data, hence we use 2006 as the starting point in

Table 6: Household consumption growth following housing crash: model vs. data.

Consumption growth, 2007-2009	High leverage	Others
Data (non-boom states)	-3.3	-2.7
Data (boom states)	-14.7	-6.8
Model (debt/house value)	-7.6	-2.7
Model (debt/assets)	-5.6	-3.9

our simulation). The high-leverage households in the model experience a sharper drop in consumption during the Great Recession than a typical household, with a cumulative decline of by 7.6% when leverage is measured as combined loan to home value ratio, and by 5.6% when leverage is measured as debt over total assets (house value plus liquid assets). The other households on average exhibit much smaller change (a drop of 2.7% or 3.9% based on the two measures of leverage).

The consumption decline experienced by the highly leveraged households in the model is comparable to that observed in the PSID for the households in the states that experienced the housing boom (where leverage grew particularly strongly).[24] In these states, households in the top quintile of leverage (measured as debt over total assets) experienced a 14.7% cumulative decline in consumption between 2007 and 2009, compared to a 6.8% decline for the other households. In contrast, households in the non-boom states only experienced a roughly 3% drop in consumption, with the decline only slightly larger for the high leverage group. This pattern is broadly consistent with evidence in Mian, Rao, and Sufi (2013) of a "debt overhang" effect, whereby households whose leverage grew the most during the boom period experienced the sharpest declines in consumption subsequently.

While deleveraging as a cause of consumption declines during the Great Recession has been a focus of much of the recent theoretical work (e.g., see Midrigan and Philippon (2011)), its mechanism is not fully understood. For example, Justiniano, Primiceri, and

[24] Dynan (2012) refers to the states that are in the top quartile of house price appreciation between 2000 and 2006 as "boom states" (they are Arizona, California, Connecticut, the District of Columbia, Florida, Hawaii, Maryland, Nevada, New Jersey, New York, Rhode Island, and Virginia).

Tambalotti (2013) argue that long-term mortgage debt that is prevalent in the data insulates households from the effect of a house price decline, hence models with one-period mortgage debt potentially overstate the magnitude of the consumption decline following the bust. Our model demonstrates that this need not be the case, as we are able to produce a substantial deleveraging effect even though we don't fully match the level of indebtedness of the most levered households. Moreover, while we allow for short-term mortgage debt (HELOC) that might need to be adjusted every period in addition to long-term mortgages, the most levered households in our model also have a nontrivial amount of positive liquid assets on average.

Can the large drops in consumption be attributed to the pure wealth effect, as argued by Kaplan, Mitman, and Violante (2017), or does leverage play a separate role? We examine the model-implied consumption growth during the Great Recession (between 2007 and 2009) across groups of households differentially affected by the increase in leverage versus declining housing wealth brought on by the house price drop. In Figure 8, we plot average consumption growth for different groups of households sorted on three different measures of exposure to the housing shock as well as the income shock: (1) leverage measured as combined loan balance relative to income (LTI), (2) combined loan-to-value (LTV), and (3) housing value relative to total household wealth (i.e. liquid assets, home equity, plus a proxy for human wealth).[25] For each measure, we plot average consumption growth in years 2008 and 2009 relative to 2007 for five different percentile groups: bottom quintile (labeled "20%"), middle quintile ("50%"), top quintile ("80%"), top decile ("90%"), the top ventile ("95%") of the distribution.

The evidence in Figure 8 indicates that the sharpest drops in consumption are experienced by the most highly levered households, those in the top 10 or 5 percentiles of loan to value (middle panel) and, especially, loan to income (left panel). The consumption drops are as large as 6% for the highest LTV percentiles, and 8 for the top ventile of LTI. These are the households that are potentially most constrained in their ability to borrow, even though some of them have substantial liquid assets accumulated during the boom years, as shown

[25] While human wealth cannot be calculated exactly, we use a simple proxy for it based on current year's labor income times a multiple of 2, which is rather conservative.

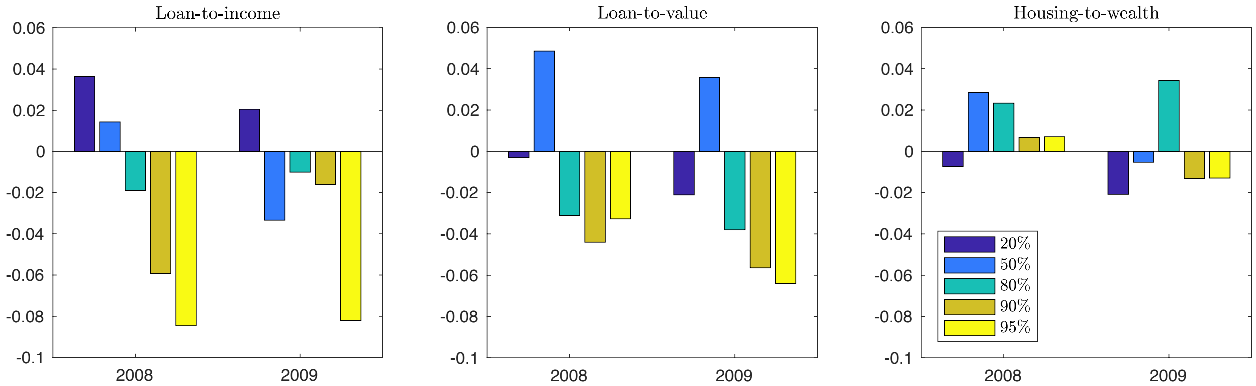

Figure 8: **Cumulative consumption growth relative to 2007.** The three panels show the cumulative consumption growth for the cross section of households sorted on LTI, LTV, and house value to total wealth in 2007.

above. However, these declines don't seem to be driven entirely by the drop in housing wealth, which is largest for the most levered households, in proportion to their home equity, which is naturally small. This is because home equity does not constitute a particularly large fraction of total wealth for these households, precisely because it is so small, relative to their liquid assets and labor income. At the same time, households for whom housing is a large component of total wealth, do not appear to suffer a large drop in consumption (at most 2%, as indicated in the right panel). Thus, according to our model, the "wealth effect" is not the primary driver of consumption decline experienced by households during the housing bust.

Why does high leverage induce large drops in consumption during the Great Recession? Despite the fact that long-term mortgages enable households to "ride out" bad times by only paying interests on the loans (as opposed to being forced to mechanically de-lever in the case on short-term mortgages), in our model, the tightening of the collateral constraints, lower expected income growth, increased uncertainty about future labor income, and costly mortgage defaults jointly generate a strong precautionary motive. In response, households reduce leverage and improve liquid asset positions, which entails cutting consumption. Note that our model might be overstating the decline in consumption among the most levered households, possibly due to the fact that it appears to produce "too little" default in the baseline specification that we employ here. Since default acts as state contingent debt, it allows for partial consumption insurance if households do not fear losing the benefits of both homeownership and access to home equity borrowing too much. However, if household see default as very costly, as our estimates suggest, they reduce leverage as a precautionary measure - this is of course a familiar mechanism from models of capital structure in the corporate finance literature.[26]

5.5 The housing boom and bust: regional variation

As discussed above, regional heterogeneity in house price fluctuations during the recent boom-bust episode might be important for understanding the wide variation in household indebtedness and the subsequent deleveraging. Indeed, Mian and Sufi (2010) document an important piece of empirical evidence in support of the effect of house prices on household borrowing. They use a measure of elasticity of housing supply developed by Saiz (2010) to show that U.S. MSAs with relatively inelastic supply of housing, which experienced fast house price growth prior to the Great Recession, saw a dramatic increase in household leverage due to home equity withdrawal, while MSAs with more elastic housing supply that had not experienced such a run-up in prices did not. In this section, we examine our model's predictions about the cross-sectional household behavior that allows for variation in the degree of the housing boom (and bust) but with otherwise similar macroeconomic conditions.

Since there is no heterogeneity in house price dynamics in our model, we approach the above evidence by conducting a counterfactual experiment. Along with our baseline model we consider two scenarios that are broadly representative of the "inelastic" and the "elastic" cases. Specifically, we solve the model using the same set of parameters as in the baseline model except those governing the stochastic process for house prices. In the "inelastic" case we let the volatility of transitory innovations to house prices be twice as large as the baseline case. This is motivated by the fact that the volatility of the transitory component of state-level

[26] Our model abstracts from other forces that might reduce the value of homeownership (and, consequently, default costs), such as exogenous moving shocks due to job transitions, etc., or access to non-collateralized borrowing, such as credit cards.

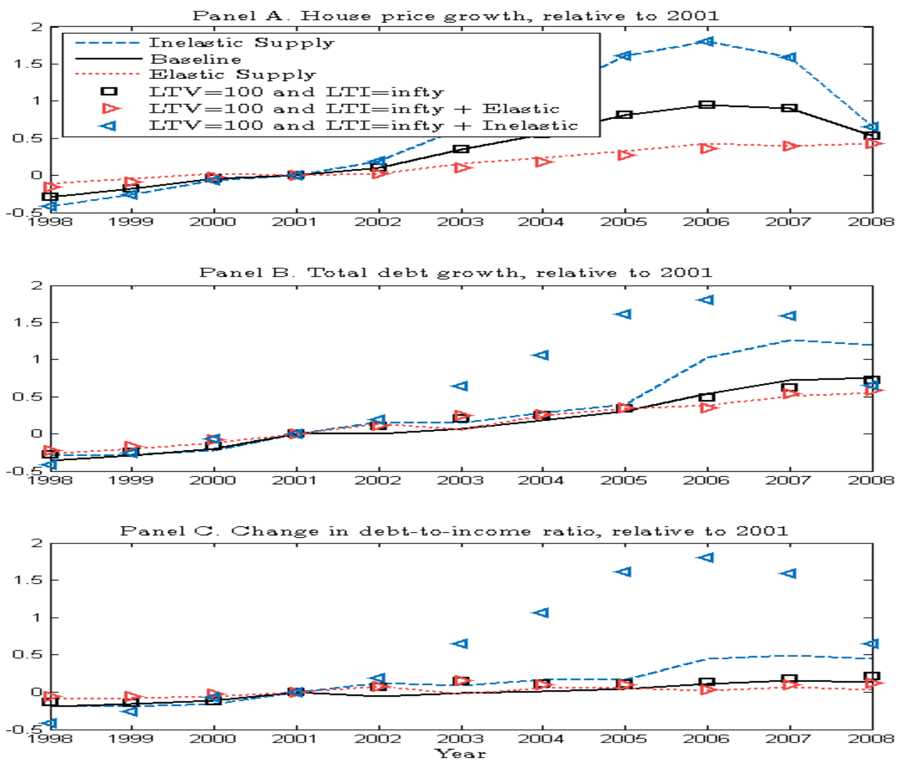

Figure 9: **Replicating Mian and Sufi (2010) evidence on household debt.**

house price indices in the "boom states" (using the definition in section 5.4 that is based on Dynan (2012)) is on average twice as high as the national average. In the "elastic" case we instead assume that the ratio of real house price to real income is constant, i.e. $p_t^H = 1$. This assumption captures the notion that in areas with elastic supply housing prices are closely aligned with construction costs (e.g., see Glaeser, Gyourko, and Saiz (2008)). Since wages are a large component of these costs, we expect house prices to be roughly proportional to labor income in the elastic areas. In addition, we perform the same experiment using the version of the model where collateral and debt service constraints are relaxed, allowing for up to 100% LTV ratio and an infinite LTI ratio, as in specification (5) of Table 5.

We plot the simulated total debt growth and changes in debt-to-income ratio over the decade 1998-2008 in Figure 9, analogous to Figure 1 in Mian and Sufi (2010). Panel A depicts the cumulative growth in house prices under the "inelastic" scenario and under the "elastic" scenario, as well as in the baseline model. The inelastic case exhibits a more rapid rise and a sharper drop in house prices than the baseline, whereas the elastic case shows only moderate growth in house prices, driven by the increase in aggregate income, consistent with the Mian-Sufi data.

Panels B and C depict the evolution of the total housing debt and the debt-to-income ratio under the two scenarios. Under the inelastic scenario with significant house price appreciation, household debt grows dramatically, especially during the period 2005-2008, both in total amount and relative to income. Compared to the Mian-Sufi data, the inelastic case overstates the total debt growth and understates the increase in debt-to-income ratio. One possible explanation for this discrepancy is that low-income households contribute more to the debt growth in the data than in our model. If so, relaxing the LTI constraint during the housing boom (while limiting the growth in total debt) will help make low-income households experience a greater increase in mortgage debt.

Indeed, we find that relaxing the borrowing constraints makes the increase in total debt and debt-to-income ratio more dramatic, especially for the latter. Both total debt and debt/income grow by 180% by 2006 and then contract by roughly three quarters of this magnitude over the subsequent two years. In contrast, under the "elastic" scenario, total debt and debt-to-income ratio stay relatively flat over the entire period, broadly in line with the evidence documented by Mian and Sufi (2010). Therefore, according to our model, the relaxation of liquidity constraints resulting from a house price run up and loosening of lending standards can jointly account for the observed increase in household leverage in a rational framework, insofar as it can be consistent with the observed path of house prices.

6 Concluding Remarks

We present an estimated structural model of household mortgage debt and liquidity management that accounts for a range of key features of both the historical time-series and the cross-sectional facts on mortgage refinancing, household leverage, and consumption. The model can be useful for quantitative evaluation of economic policies aimed at supporting household balance sheets via the mortgage market.

Our simulation-based evidence also demonstrates that the interaction between interest rates and household liquidity constraints is important for assessing the effect of monetary policy on refinancing activity. When many households are liquidity constrained, their refinancing behavior becomes insensitive to changes in interest rates, especially in the face of depressed values of housing collateral or high debt service ratios (this is further corroborated by the recent evidence in Beraja, Fuster, Hurst, and Vavra (2017) and Di Maggio, Kermani, and Palmer (2016)). At the same time, our analysis suggests that a monetary easing in the early stages of an economic downturn, when both aggregate income falls and its cross-sectional dispersion rises, elicits stronger refinancing activities than what standard models would predict based solely on interest rate changes, unless it is accompanied with an additional tightening of lending standards.

However, our estimated model overstates the quantitative magnitude of the effect of liquidity demand on mortgage refinancing. It overshoots the average refinancing rate (11% vs. 7% in the data) and the size of cash-outs conditional on refinancing (by nearly a factor of 3) while understating the part of refinancing for rate reasons, suggesting that the estimated costs of refinancing are likely too low and the "house as ATM" channel might be too strong in the model. The model under-predicts the leverage and refinancing rate for high-income households (they are more debt-dependent in the data). It also under-predicts the size of cash-outs for high-leverage households, suggesting that the borrowing constraint was less binding in the data at least over some of the sample period than in the model. Relative to the data, the model generates a "delayed" refinancing boom for high-leverage households (it occurs in 2005 in our simulations instead of 2003 as in the data). This is because these households are prevented from refinancing in 2003 due to a binding borrowing constraint (subsequently relaxed by the rapid house price appreciation). Recent evidence indicates that credit constraints loosened substantially in the data in 2003, potentially explaining this discrepancy (e.g., Justiniano, Primiceri, and Tambalotti (2017)).

The partial equilibrium nature of our model prevents us from fully capturing the implications of refinancing-related frictions on aggregate consumption, which might include feedback effects into house prices (e.g. due to the collateral value of housing or to default-driven fire sales), mortgage rates (due to fluctuations in prepayment and default risk, as well as cross-sectional heterogeneity), and, ultimately, aggregate household income (due to balance-sheet-driven shifts in consumer demand). Exploring these links in general equilibrium settings is an exciting challenge and a subject of ongoing research.

References

Alvarez, Fernando E., Luigi Guiso, and Francesco Lippi, 2010, Durable consumption and asset management with transaction and observation costs, Nber working paper no. 15835.

Attanasio, Orazio, Andrew Leicester, and Matthew Wakefield, 2011, Do house prices drive consumption growth? The coincident cycles of house prices and consumption in the UK, Journal of the European Economic Association 9, 399–435.

Attanasio, Orazio P., and Guglielmo Weber, 1995, Is consumption growth consistent with intertemporal optimization? evidence from the consumer expenditure survey, Journal of Political Economy 103, 1121–1157.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, Risks for the long run: Estimation with time aggregation, Working Paper 18305 National Bureau of Economic Research.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra, 2017, Regional heterogeneity and monetary policy, Working Paper 23270 National Bureau of Economic Research.

Bertola, Giuseppe, Luigi Guiso, and Luigi Pistaferri, 2005, Uncertainty and consumer durables adjustment, The Review of Economic Studies 72, 973–1007.

Brav, Alon, George M. Constantinides, and Christoper Geczy, 2002, Asset pricing with heterogeneous consumers and limited participation: Empirical evidence, Journal of Political Economy 110, 793–824.

Calomiris, Charles W, Stanley D Longhofer, and William Miles, 2012, The housing wealth effect: The crucial roles of demographics, wealth distribution and wealth shares, Discussion paper National Bureau of Economic Research.

Campbell, Jeffrey C., and Zvi Hercowitz, 2005, The role of collateralized household debt in macroeconomic stabilization, NBER Working Paper Series No. 11330.

Campbell, John Y., and Joao F. Cocco, 2003, Household risk management and optimal mortgage choice, The Quarterly Journal of Economics 118, 1449–1494.

——— , 2007, How do house prices affect consumption? evidence from micro data, Journal of Monetary Economics 54, 591–621.

——— , 2015, A model of mortgage default, Journal of Finance 70, 1495–1554.

Caplin, Andrew, Charles Freeman, and Joseph Tracy, 1997, Collateral damage: How refinancing constraints exacerbate regional recessions, Journal of Money, Credit and Banking 29, 496–516.

Carroll, Christopher, 2001, Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation), The B.E. Journal of Macroeconomics 1, 1–38.

——— , Jirí Slacálek, and Martin Sommer, 2012, Dissecting saving dynamics: Measuring credit, uncertainty, and wealth effects, JHU Working Paper.

Carroll, Christopher D., 2000, Why do the rich save so much?, in Joel Slemrod, ed.: Does Atlas Shrug? The Economic Consequences of Taxing the Rich (Harvard University Press: Cambridge, MA).

———, Misuzu Otsuka, and Jiri Slacalek, 2011, How large are housing and financial wealth effects? A new approach, Journal of Money, Credit and Banking 43, 55–79.

Case, Karl E., John M. Quigley, and Robert J. Shiller, 2011, Wealth effects revisited 1978-2009, NBER Working Paper Series No. 16848 Yale University.

Chatterjee, Satyajit, and Burcu Eyigungor, 2015, A quantitative analysis of the u.s. housing and mortgage markets and the foreclosure crisis, Review of Economic Dynamics 18, 165 – 184.

Chen, Hui, Jianjun Miao, and Neng Wang, 2010, Entrepreneurial finance and nondiversifiable risk, Review of Financial Studies 23, 4348–4388.

Corbae, Dean, and Erwan Quintin, 2013, Leverage and the foreclosure crisis, NBER Working Paper.

Davis, Morris A., and Francois Ortalo-Magne, 2011, Household expenditures, wages, rents, Review of Economic Dynamics 14, 248 – 261.

DeNardi, Mariacristina, 2004, Wealth inequality and intergenerational links, Review of Economic Studies 71, 743–768.

Di Maggio, Marco, Amir Kermani, and Christopher Palmer, 2016, How quantitative easing works: Evidence on the refinancing channel, Working Paper 22638 National Bureau of Economic Research.

Dridi, Ramdan, Alain Guay, and Eric Renault, 2007, Indirect inference and calibration of dynamic stochastic general equilibrium models, Journal of Econometrics 136, 397 – 430.

Duffie, Darrell, and Kenneth J Singleton, 1993, Simulated moments estimation of Markov models of asset prices, Econometrica 61, 929–52.

Dynan, Karen, 2012, Is a household debt overhang holding back consumption?, Brookings Papers on Economic Activity pp. 299–362.

Epstein, Larry, and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption growth and asset returns I: A theoretical framework, Econometrica 57, 937–969.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh, 2011, The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium, NBER Working Paper Series No. 15988.

Fernandez-Villaverde, Jesus, and Dirk Krueger, 2011, Consumption and saving over the life cycle: How important are consumer durables?, Macroeconomic Dynamics 15, 725–770.

Gallant, A. Ronald, and George E. Tauchen, 1996, Which moments to match, Econometric Theory 16, 657–681.

Gerardi, Kristopher, Kyle F. Herkenhoff, Lee E. Ohanian, and Paul S. Willen, 2018, Cant pay or wont pay? unemployment, negative equity, and strategic default, The Review of Financial Studies 31, 1098–1131.

Glaeser, Edward L., Joseph Gyourko, and Albert Saiz, 2008, Housing supply and housing bubbles, Journal of Urban Economics 64, 198 – 217.

Gomes, Francisco, and Alexander Michaelides, 2005, Optimal life-cycle asset allocation: Understanding the empirical evidence, Journal of Finance 60, 869–904.

Gourieroux, C, A Monfort, and E Renault, 1993, Indirect inference, Journal of Applied Econometrics 8, S85–118.

Gourinchas, Pierre-Olivier, and Jonathan A. Parker, 2002, Consumption over the life cycle, Econometrica 70, 47–89.

Greenwald, Daniel, 2017, The mortgage credit channel of macroeconomic transmission, .

Gross, David B., and Nicholas S. Souleles, 2002, Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data, The Quarterly Journal of Economics 117, 149–185.

Guerrieri, Veronica, and Guido Lorenzoni, 2011, Credit crises, precautionary savings, and the liquidity trap, NBER Working Papers 17583 National Bureau of Economic Research, Inc.

Guiso, Luigi, and Paolo Sodini, 2013, Household finance: An emerging field, in Milton Harris George M. Constantinides, and Rene M. Stulz, ed.: Handbook of the Economics of Finance SET, vol. 2, Part B of Handbook of the Economics of Finance . pp. 1397 – 1532 (Elsevier).

Guvenen, Fatih, Serdar Ozkan, and Jae Song, 2012, The nature of countercyclical income risk, NBER Working Papers 18035 National Bureau of Economic Research, Inc.

Hall, Robert E., 1988, Intertemporal substitution and consumption, Journal of Political Economy 96, 339–357.

Hansen, Lars P., 1982, Large sample properties of generalized method of moments estimators, Econometrica 50, 1029–1054.

Hansen, Lars Peter, John Heaton, Junghoon Lee, and Nikolai Roussanov, 2007, Intertemporal substitution and risk aversion, , vol. 6, Part 1 of Handbook of Econometrics . pp. 3967 – 4056 (Elsevier).

He, Chao, Randall Wright, and Yu Zhu, 2012, Housing and liquidity, Working Paper, University of Wisconsin.

Hennessy, Christopher A., and Toni M. Whited, 2005, Debt dynamics, Journal of Finance 60, 1129–1165.

Hurst, Erik, Benjamin Keys, Amit Seru, and Joseph Vavra, 2014, Regional redistribution through the U.S. mortgage market, Working paper.

Hurst, Erik, and Frank Stafford, 2004, Home is where the equity is: Mortgage refinancing and household consumption, Journal of Money, Credit and Banking 36, 985–1014.

Iacoviello, Matteo, and Marina Pavan, 2013, Housing and debt over the life cycle and over the business cycle, Journal of Monetary Economics 60, 221 – 238.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti, 2013, Household leveraging and deleveraging, NBER Working Papers 18941 National Bureau of Economic Research, Inc.

Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti, 2017, The mortgage rate conundrum, Working Paper 23784 National Bureau of Economic Research.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante, 2017, The housing boom and bust: Model meets evidence, Working Paper 23694 National Bureau of Economic Research.

Kaplan, Greg, and Giovanni L. Violante, 2011, A model of the consumption response to fiscal stimulus payments, NBER Working Papers 17338 National Bureau of Economic Research.

Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, 2010, Did securitization lead to lax screening? Evidence from subprime loans, Quarterly Journal of Economics 125, 307–362.

Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vikrant Vig, 2012, Mortgage Financing in the Housing Boom and Bust . pp. 143–204 (University of Chicago Press).

Kiyotaki, Nobuhiro, Alexander Michaelides, and Kalin Nikolov, 2011, Winners and losers in housing markets, Journal of Money, Credit and Banking 43, 255–296.

Koijen, Ralph S.J., Otto Van Hemert, and Stijn Van Nieuwerburgh, 2009, Mortgage timing, Journal of Financial Economics 93, 292–324.

Laibson, David, 1997, Golden eggs and hyperbolic discounting, Quarterly Journal of Economics 112, 443–477.

——— , and Johanna Mollerstrom, 2010, Capital flows, consumption booms and asset bubbles: A behavioural alternative to the savings glut hypothesis, The Economic Journal 120, 354–374.

Laibson, David, Andrea Repetto, and Jeremy Tobacman, 2003, A debt puzzle, in Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Strother Phelps . p. 228 (Princeton University Press).

——— , 2007, Estimating discount functions with consumption choices over the lifecycle, NBER Working Paper.

Landvoigt, Tim, 2017, Housing demand during the boom: The role of expectations and credit constraints, The Review of Financial Studies 30, 1865–1902.

——— , Monika Piazzesi, and Martin Schneider, 2012, The housing market(s) of San Diego, Working Paper 17723 National Bureau of Economic Research.

Laufer, Steven, 2017, Equity extraction and mortgage default, Review of Economic Dynamics.

Lustig, Hanno, and Stijn Van Nieuwerburgh, 2010, How much does household collateral constrain regional risk sharing?, Review of Economic Dynamics 13, 265–294.

Meghir, Costas, and Luigi Pistaferri, 2004, Income variance dynamics and heterogeneity, Econometrica 72, 1–32.

Mian, Atif R., Kamaleshand Rao, and Amir Sufi, 2013, Household balance sheets, consumption, and the economic slump, Chicago Booth Working Paper.

Mian, Atif R., and Amir Sufi, 2010, House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis, American Economic Review, forthcoming, avaliable at http://ideas.repec.org/p/nbr/nberwo/15283.html.

Midrigan, Virgiliu, and Thomas Philippon, 2011, Household leverage and the recession, Working Paper 16965 National Bureau of Economic Research.

Ortalo-Magné, François, and Sven Rady, 2006, Housing market dynamics: On the contribution of income shocks and credit constraints, The Review of Economic Studies 73, 459–485.

Parker, Jonathan A., and Annette Vissing-Jørgensen, 2009, Who bears aggregate fluctuations and how?, American Economic Review 99, 399–405.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel, 2007, Housing, consumption and asset pricing, Journal of Financial Economics 83, 531–569.

Rios-Rull, Jose-Victor, and Virginia Sanchez-Marcos, 2008, An aggregate economy with different size houses, Journal of the European Economic Association 6, 705–714.

Roussanov, Nikolai, 2010, Diversification and its discontents: Idiosyncratic and entrepreneurial risk in the quest for social status, Journal of Finance 65, 1755–1788.

Saiz, Albert, 2010, The geographic determinants of housing supply, The Quarterly Journal of Economics 125, 1253–1296.

Smith, A A, Jr, 1993, Estimating nonlinear time-series models using simulated vector autoregressions, Journal of Applied Econometrics 8, S63–84.

Storesletten, Kjetil, Chris Telmer, and Amir Yaron, 2004, Cyclical dynamics in idiosyncratic labor market risk, Journal of Political Economy 112, 695–717.

———, 2007, Asset pricing with idiosyncratic risk and overlapping generations, Review of Economic Dynamics 10, 519–548.

Tauchen, George, and Robert Hussey, 1991, Quadrature based methods for obtaining approximate solutions to nonlinear asset pricing models, Econometrica 59, 371–396.

Telyukova, Irina A., 2013, Household need for liquidity and the credit card debt puzzle, Review of Economic Studies, forthcoming.

Vissing-Jørgensen, Annette, 2002, Limited stock market participation and the elasticity of intertemporal substitution, Journal of Political Economy 110, 825–853.

———, 2007, Household finance: The liability side, Focus Session Introduction, Gerzensee European Summer Symposium on Financial Markets.

———, and Orazio P. Attanasio, 2003, Stock-market participation, intertemporal substitution, and risk-aversion, American Economic Review 93, 383–391.

Wachter, Jessica A., and Motohiro Yogo, 2010, Why do household portfolio shares rise in wealth?, Review of Financial Studies 23, 3929–3965.

Weil, Philippe, 1990, Non-expected utility in macroeconomics, Quarterly Journal of Economics 105, 29–42.

Woodward, Susan E., and Robert E. Hall, 2010, Consumer confusion in the mortgage market: Evidence of less than a perfectly transparent and competitive market, American Economic Review 100, 511–15.

Internet Appendix

A State level evidence on counter-cyclical refinancing

To investigate the response of mortgage refinancing to economic activity further, we use data on the origination of home mortgage loans at the state level. This potentially allows us to separate the effect of low interest rates from that of deteriorating economic conditions, insofar as the local economic activity variables are less synchronized with the interest rates than are aggregate quantities, and that households cannot diversify away state-level shocks.[27]

We use quarterly data on the mortgage loans (both refinance and purchase) for each of the 50 states and D.C., based on aggregated Home Mortgage Disclosure Act (HMDA) reporting. We regress the quarterly changes in the number of loans taken in order to refinance existing mortgages (adjusted by the state population) on measures of economic conditions. We use three such measures, specifically growth rates of nonfarm payroll employment, of the State Coincident Economic Activity Index ($CEAI$), which combines information contained in nonfarm payrolls, unemployment, hours worked and wages, and trends with the Gross State Product (GSP), and of the total personal income (TPI), deflated using the national consumer price index.[28] We use year-on-year (log) growth rates of quarterly levels of these measures as the main explanatory variables.

House prices determine both the motive to refinance due to a wealth effect and the ability of households to borrow against the value of their homes (perhaps for reasons unrelated to consumption smoothing). Since economic conditions are correlated with the level of house prices, refinancing activity could be high under good economic conditions due to high house prices. Thus, to better capture the effect of consumption smoothing on refinancing, it is important to control for house price appreciation in our regression. We use the FHFA house price indices for the 50 states and DC as our measure of house prices. As before, we also control for aggregate variables: the 30 year mortgage rate (contemporaneous and lagged by one year) and the short-term interest rate.

We run pooled time series/cross-sectional regressions of the form:

$$\begin{aligned} REFI_t^{State} = {} & \beta_{Cycle}^{REFI} Cycle_t^{State} + \beta_H^{REFI} \Delta HPI_t^{State} + \beta_{CH}^{REFI} Cycle_t^{State} \times HPI_t^{State} + \bar{R}_t^i \\ & + \beta_w^{REFI} WAC_t^{State} + \beta_r^{REFI} R_t^{3M} + \beta_R^{REFI} R_t^{M30} + \beta_{Rl}^{REFI} R_{t-4}^{M30} + \beta_t + \beta_{State} + \epsilon_t, \end{aligned} \tag{A.1}$$

where $REFI_t^{State}$ is the number of refinance loans originated in state i over the quarter t, scaled by the state's population in the prior year.[29] $Cycle^{State}$ is the variable that measures state-level aggregate economic conditions, ΔHPI_t measures house price appreciation using the 2-year growth in the FHFA state-level house price index that captures appreciation of the mortgaged properties, $\bar{R}_t^i$ is the average rate on newly originated conventional mortgages

[27] Hurst, Keys, Seru, and Vavra (2014) show that there is essentially no cross-state variation in mortgage rates on loans originated by government-sponsored enterprises (GSEs).

[28] Unlike the payroll employment and personal income measures, $CEAI$ is not available for D.C.

[29] We obtain similar results using refinance loan volume scaled by total personal income in the state.

in state i over the past year (also provided by FHFA),[30] WAC_t^{State} is the weighted average coupon on conforming mortgage loans outstanding in the state in the first month of the quarter that summarizes the rates currently paid by borrowers, $\mathbf{b}_t$ is the vector of quarter fixed effects that captures aggregate information not contained in other variables, and $\mathbf{b}_{State}$ a vector of state fixed effects. State fixed effects are important since there is substantial heterogeneity across states in the fixed costs associated with refinancing a mortgage (such as title insurance, taxes, etc.), which result in different average levels of refinancing as well as its sensitivity to aggregate variables. Given this specification, we are identifying the effect of within-state variation in economic conditions on refinancing. We include the lagged $Cycle$ variable to capture delayed response of households to economic conditions, and include an interaction term between $Cycle$ and the house price growth, orthogonalized with respect to both variables, to test whether higher level of house prices help relax the borrowing constraint especially in bad times.

Table A.1 presents the results of the state-level regressions for different specifications (two different economic activity measures). The coefficients on the state-level business cycle variables in the first column are all negative and statistically significant in all but one specification (TPI without time fixed effects), consistent with the view that households are more likely to refinance their mortgages in a downturn. The state-level cycle variable remains significantly negatively related to refinancing when the quarter fixed effects are included, indicating that their presence does not simply proxy for variation in the aggregate term structure variables.

As expected, house price appreciation is positively related to refinancing. In fact, the effects of the business cycle variables become stronger (more negative) after house price appreciation is taken into account, which helps tease out the rise in refinancing in good times due to house value appreciation (results without house price index are not reported). Moreover, the interaction terms of house prices and the cycle variables are negative and typically statistically significant, suggesting that higher levels of house prices are particularly important for refinancing during economic downturns.

Both the 30-year mortgage rates and the short-term interest rate have a significant negative effect on refinancing, as expected. Similarly, the WAC has a significant positive coefficient, consistent with the fact that it captures the rates currently paid by borrowers, so that higher WAC translated into a greater incentive to refinance if current rates are low. In the specification with time fixed effects (where aggregate interest rates are not included) WAC has a negative coefficient, potentially due to the fact that it may capture persistent state-specific variation in mortgage spreads that we cannot control for separately without detailed state-level data on mortgage rates. Interestingly, the relationship of refinancing with contemporaneous state-level mortgage rates is positive rather than negative, although not significant with time fixed effect, suggesting that it is capturing mostly aggregate variation in mortgage spreads, which are positively related to both default and prepayment risk, and are likely to increase with the rising demand for mortgage loans in a particular state.

[30] This variable is reported at annual frequency; we generate quarterly observations via linear interpolation.

Table A.1: State-level refinancing activity

	$Cycle_t$	HPI_t	$C_t \times H_t$	WAC	$\bar{R}_t^i$	R_t^{M30}	R_t^{3M}	R_{t-4}^{M30}	$\bar{R}^2$
1	-0.29	0.17	-1.85	0.62	1.50	-1.70	-0.75	-0.20	0.61
Robust	(0.05)	(0.01)	(0.51)	(0.03)	(0.22)	(0.11)	(0.06)	(0.11)	
NW	(0.05)	(0.01)	(0.39)	(0.05)	(0.22)	(0.12)	(0.06)	(0.12)	
2	-0.24	0.10	-0.64	-2.74	0.32				0.89
Robust	(0.05)	(0.01)	(0.27)	(0.70)	(0.41)				
NW	(0.05)	(0.01)	(0.20)	(0.67)	(0.37)				
3	-0.10	0.16	-1.29	0.64	1.56	-1.79	-0.80	-0.23	0.60
Robust	(0.03)	(0.01)	(0.42)	(0.04)	(0.24)	(0.12)	(0.06)	(0.11)	
NW	(0.03)	(0.01)	(0.34)	(0.05)	(0.23)	(0.12)	(0.07)	(0.12)	
4	-0.14	0.10	-0.47	-2.62	0.36				0.89
Robust	(0.04)	(0.01)	(0.19)	(0.70)	(0.42)				
NW	(0.03)	(0.01)	(0.13)	(0.69)	(0.37)				
5	0.01	0.15	-1.89	0.61	1.84	-1.89	-1.00	-0.32	0.60
Robust	(0.03)	(0.01)	(0.54)	(0.04)	(0.27)	(0.14)	(0.06)	(0.11)	
NW	(0.03)	(0.01)	(0.37)	(0.05)	(0.26)	(0.13)	(0.07)	(0.13)	
6	-0.10	0.09	-0.36	-2.63	0.18				0.89
Robust	(0.03)	(0.01)	(0.25)	(0.70)	(0.44)				
NW	(0.03)	(0.01)	(0.22)	(0.70)	(0.39)				

NOTE: Quarterly data, 1993.III - 2009.IV (time subscript t is in monthly units). The dependent variable is the total number of newly originated refinance loans in the state over a quarter relative to the rescaled population of the state for the previous year (based on HMDA data). Cycle refers to the year-on-year growth in either the non-farm payroll employment index scaled by the state population (*Payroll*, specifications 1 - 2), State Coincident Economic Activity index in columns (*CEAI*, specifications 3 - 4), or the Total Personal Income (*TPI*, deflated using the CPI, specifications 5 - 6). HPI is the two-year growth rate of the state-level house price index. $C_t \times H_t$ is the orthogonalized interaction term, i.e. the residual from regressing the product of Cycle and HPI on a constant and both of these variables. WAC is weighted average coupon rate for conforming fixed-rate mortgages (equal-weighted average across FNMA and FHLMC loans) in a given state. $\bar{R}_t^i$ is the average coupon rate on all newly-originated conventional prime loans in the state over the quarter. Specifications 2, 4 and 6 have quarter fixed effects. Standard errors are in brackets (Robust are clustered by state, and NW are Newey-West with 20 lags).

B Household problem

In this section, we specify the problem for homeowners and renters. In order to simplify notation, we drop subscripts t and use primes to denote next period variables.

Homeowner problem The problem for homeowner i is to choose real non-housing consumption c_i, house size h_i, the position in the liquid asset (or HELOC) a'_i, as well as the decision to refinance or repay early (both of which result in a new mortgage balance b'_i), sell the house, or default on the debt, so as to maximize the expected lifetime utility of real consumption. Denoting the refinancing decision by the indicator $I^{RF}_{i,t}$ (with $I^{RF}_{i,t} = 1$ for refinancing at time t and 0 otherwise), the dynamics of the mortgage rate $k_{i,t}$ can be written as

$$k_{i,t+1} = k_{i,t}\ (1 - I^{RF}_{i,t}) + R_t I^{RF}_{i,t}. \tag{A.2}$$

The household problem can be formalized as follows,

$$U^h_i(a_i, b_i, k_i, h_i, v_i) = \max_{a'_i, b'_i, h'_i, I^{RF}_i} \left[(1-\delta) \left((h_i Y)^{\nu}\, c_i^{1-\nu} \right)^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E} \left[\max \left(U^{h'}_i, U^{hr'}_i, U^{hd'}_i \right)^{1-\gamma} \right]^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}} \tag{A.3}$$

subject to a

$$\begin{aligned} c_i P + \frac{a^{+\prime}_i}{1 + (1-\tau) r} + \frac{a^{-\prime}_i}{1 + r^{HL}} + b_i = (1-\tau)(y_i - k_i b_i) + a_i + b'_i - \phi(b'_i; V)\ I^{RF}_i \\ + I^M_i P^H \left((1-\phi_h)\, h_i - (1+\phi_h)\, h'_i \right) \qquad \text{(A.4)} \\ (b'_i - b_i)\ (1 - I^{RF}_i) \leq 0, \\ c_i, b'_i \geq 0, \end{aligned}$$

along with the law of motion for mortgage rate k_i (A.2), the LTV and LTI constraints (8) and (9), and the upper bound on HELOC (10). We denote the value function of the household in the homeowner state by $U^h_i(a_i, b_i, k_i, h_i, v_i)$, by $U^{hr}_i(a_i, b_i, k_i, h_i, v_i)$ in a state of transition from homeowner to renter by selling the home, and by $U^{hd}_i(a_i, b_i, k_i, h_i, v_i)$ in a state of transition from homeowner to renter by defaulting on the mortgage.

Upon transition from homeownership to renter state the proceeds from selling the house $(1-\phi_h) h_i P^H$ are added to the resource constraint while the mortgage and HELOC borrowing must be repaid. The problem for the household making the transition from the homeowner to the renter state by selling its home is then given by

$$U^{hr}_i(a_i, b_i, k_i, h_i, v_i) = \max_{a'_i} \left[(1-\delta) \left((h_i Y)^{\nu}\, c_i^{1-\nu} \right)^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E} \left[U^r_i(a'_i, v'_i)^{1-\gamma} \right]^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \tag{A.5}$$

subject to

$$\begin{aligned} c_i P + \frac{a'_i}{1 + (1-\tau) r} &= (1-\tau)(y_i - k_i b_i) + a_i + (1-\phi_h) h_i P^H - b_i, \qquad \text{(A.6)} \\ a'_i, c_i &\geq 0, \end{aligned}$$

where $U_i^r(a_i, v_i)$ denotes the value function of an unrestricted renter who is allowed to buy a house immediately.

If a household defaults on its mortgage, it also becomes a renter, but with the added restriction that it will be excluded from the housing market for a period of time. This transition problem is given by

$$U_i^{hd}(a_i, b_i, k_i, h_i, v_i) = \max_{a_i'} \left[(1-\delta) \left((h_i Y)^\nu c_i^{1-\nu} \right)^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E} \left[U_i^d(a_i', v_i')^{1-\gamma} \right]^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}} \tag{A.7}$$

subject to

$$c_i P + \frac{a_i'}{1+(1-\tau)r} = (1-\tau)y_i + \zeta a_i^+, \tag{A.8}$$
$$a_i', c_i \geq 0,$$

where $U_i^d(a_i, v_i)$ denotes the value function of a restricted renter who is currently excluded from the housing market due to defaulting on its mortgage. In both (A.7) and (A.8), the constraint $a_i \geq 0$ is due to the fact that HELOC is unavailable to renters.

Renter problem For convenience, we define three different types of renters: unrestricted renter, restricted renter, and a renter in transition to become a homeowner, with value functions $U_i^r(a_i, v_i)$, $U_i^d(a_i, v_i)$, and $U_i^{rh}(a_i, v_i)$, respectively. The problem for an unrestricted renter is to choose the size of the rental house h_i^r, the non-housing consumption c_i, and the liquid assets for the next period a', such that

$$U_i^r(a_i, v_i) = \max_{h_i^r, a_i' \geq 0} \left[(1-\delta) \left((h_i^r Y)^\nu c_i^{1-\nu} \right)^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E} \left[\max \left(U_i^{rh}(a_i', v_i'), U_i^r(a_i', v_i') \right)^{1-\gamma} \right]^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}} \tag{A.9}$$

subject to the positivity of consumption and the budget constraint:

$$h_i^r \varpi P Y + c_i P = (1-\tau)y_i + a_i - \frac{a_i'}{1+(1-\tau)r}. \tag{A.10}$$

The intra-temporal optimization implies

$$\frac{h_i^r \varpi P Y}{c_i P} = \frac{\nu}{1-\nu}.$$

That is, the ratio of rental expense and non-housing consumption is constant. This condition helps simplify the Bellman equation (A.9) and the renter budget constraint (A.10) into

$$U_i^r(a_i, v_i) = \max_{a_i' \geq 0} \left[(1-\delta) \left(\overline{\eta} c_i \right)^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E} \left[\max \left(U_i^{rh}(a_i', v_i'), U_i^r(a_i', v_i') \right)^{1-\gamma} \right]^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}} \tag{A.11}$$

and

$$\frac{c_i P}{1-\nu} = (1-\tau)y_i + a_i - \frac{a_i'}{1+(1-\tau)r}, \tag{A.12}$$

where

$$\overline{\eta} \doteq \left(\frac{\nu}{(1-\nu)\varpi}\right)^{\nu}. \tag{A.13}$$

The transition problem for the household from the renter to the homeowner state is given by

$$U_i^{rh}(a_i, v_i) = \max_{a_i', b_i', h_i'} \left[(1-\delta)\left(\overline{\eta} c_i\right)^{\frac{1-\gamma}{\theta}} + \delta\, \mathbb{E}\left[U_i^h(a_i', b_i', k_i, h_i, s_i')^{1-\gamma}\right]^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}}, \tag{A.14}$$

subject to

$$\frac{c_i P}{1-\nu} = (1-\tau)y_i + a_i - \frac{a_i'}{1+(1-\tau)r} + b_i' - \phi(b_i'; V) - (1+\phi_h)h_i' P^H, \tag{A.15}$$
$$c_i, b_i' \geq 0,$$

the LTV and LTI constraints (8) and (9), and the constraint on HELOC (10).

The problem of a restricted (post-default) renter is given by

$$\begin{aligned} U_i^d(a_i, v_i) \quad = \quad \max_{a_i' \geq 0} \Bigg[&(1-\delta)\left(\overline{\eta} c_i\right)^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E}\left[(1-\omega)\left(U_i^d(a_i', v_i')\right)^{1-\gamma}\right]^{\frac{1}{\theta}} \\ &+ \delta \mathbb{E}\left[\omega \max\left(U_i^{rh}(a_i', v_i'), U_i^r(a_i', v_i')\right)^{1-\gamma}\right]^{\frac{1}{\theta}} \Bigg]^{\frac{\theta}{1-\gamma}}, \end{aligned} \tag{A.16}$$

subject to the positivity of consumption as well as the renter budget constraint (A.12).

Since households have homothetic preferences, we rescale the problem with respect to the price level P_t and the permanent aggregate income Y_t in order to make it stationary.

C Computation and Estimation

Prespecified parameters The parameters controlling the dynamics of the exogenous state variables as well as describing the institutional features of the model environment are summarized as

$$\Theta_0 \equiv (\mu_S, \Phi_S, \Sigma_S, \pi, \mu_y, \rho_y, \sigma_y(\cdot), \bar{H}, \tau, \kappa_0, \kappa_1, \kappa_2, \xi_{LTI}, \xi_{LTV}, \underline{a}, \zeta, \omega, \vartheta).$$

Numerical Implementation The household problem is solved numerically using a standard value function iteration (VFI) procedure on a very large grid (more than 1.9 million total grid points, with 1920 points for the exogenous states and 960 points for the endogenous states). Moreover, we need to solve the model repeatedly in the estimation. These requirements make the computational problem rather challenging. To make the estimation feasible, we programmed the numerical solution in CUDA language and ran the VFI on a Nvidia C2050 (Fermi) graphics card (with 448 CUDA cores). Since the objective function is highly nonlinear, we use a global search algorithm to ensure that the resulting estimates are not due to local minima. The estimation was implemented with a global optimization routine capable of using up to 8 graphics cards simultaneously. This (software and hardware) implementation yields a significant improvement in speed, allowing us to estimate the model in less than one week. The same estimation problem will take 400 times as long on a standard desktop computer.

Simulation-based inference In order to be able to evaluate the statistical significance of the mismatch between the target and simulated moments, as well as the uncertainty about the estimated parameter values, we need to estimate the variance-covariance matrix of the sample moments, Ξ. Since we use a combination of time-series and cross-sectional moments, using data directly is not feasible. Instead, we construct the variance-covariance matrix of the simulated moments under the null that the model is true (with the parameters set at the estimated values). In order to estimate this matrix we simulate $N_A = 80$ paths of aggregate variables and generate a panel of $N = 1000$ households using these aggregate shocks and simulated idiosyncratic shocks so that it matches the small sample length $T_D = 25$ years available in the data. For each of the aggregate paths we compute the full set of moments, and estimate the variance-covariance matrix of these moment vectors across simulations. While the simulated moments used in estimation are based on long samples of length T, i.e. are essentially population moments, the variance-covariance matrix estimated using the short-sample simulated moments measures the sampling uncertainty about the moments estimated in the data under the null of the model.

In addition, we construct standard errors for the estimated parameters from the Ξ matrix using the standard delta method,

$$var(\hat{\Theta}) = \frac{1}{T_D}(d'Wd)^{-1}d'W\Xi Wd(d'Wd)^{-1},$$

where the derivatives of the moments with respect to the parameters $d = \frac{\partial m(\Theta_0,\Theta)}{\Theta}$ are approximated using numerical finite differences.

D Data

D.1 Aggregates

Aggregate consumption, personal income, and gross domestic product are from the U.S. National Income and Product Accounts; house price index is from $S\&P$/Case-Shiller; one-year Treasury Bill rate from FRED; 30-year fixed mortgage rate is from Freddie Mac Primary Mortgage Market Survey (PMMS). Homeownership rate is from the U.S. Census (average over the time period 1990-2010 is 66.54%). The number and volume of mortgage refinancing originations, as well as the average ratio of the loan amount to income, by state, per quarter, is based on the Home Mortgage Disclosure Act (HMDA) reporting for the time period 1993-2009. Total dollar cash-out relative to total dollar refinancing volume for prime, conventional loans, as well as the fraction of loans that involve cash-out and the median ratio of new to old rate are from Freddie Mac for the time period 1993-2010.

The target regression coefficients are based on the auxiliary model. For total refinancing we estimate (1) assuming that $Z = \Delta IP_t$. For cash-out, we estimate (2) for cash-out refinancing and assume that $Z_t = \Delta PI_t$. For both cases we use the full specification estimates displayed in Table 1.

D.2 Survey of Consumer Finances

We use the SCF public data set available from the Federal Reserve Board of Governors for the years 1989, 1992, 1995, 1998, 2001, 2004, 2007, and 2010 for model estimation and evaluation. The survey is representative of the U.S. population and is designed to oversample the wealthy households. Each household is represented in the data set by 5 replicates (implicates) constructed in order to compensate for omitted information about households assets, etc.We use sampling weights provided by the SCF to allow aggregation to population totals.

The survey contains detailed information on household demographics, income, debt, and asset holdings. We define liquid assets in the SCF data as the total value of checking/savings accounts, bonds, and public equity holdings, including both directly-held stocks and mutual funds. Kaplan and Violante (2011) use a similar definition. For mortgage debt we use the first lien loan collateralized by the primary residence of the household, whereas the combined balance of all of the junior lien loans on the same residence (including second/third mortgages and home equity lines of credit) is classified as HEL(OC). Income is total family income in the calendar year (prior to the survey year); we drop households with income less than \$1000/year. House value is based on the total value of the primary residence (for homeowners). Refinancing statistics are constructed based on mortgages that are identified as refinance loans originated during the year of the survey or the prior year.

For each year we remove the top quartile of asset holdings and then pool the data across years in computing averages and distributions.

D.3 Panel Study of Income Dynamics

We use Panel Study of Income Dynamics waves for the years 1999, 2001, 2003, 2005, 2007, 2009, 2011, and 2013.

The main survey contains detailed information on household demographics, income, and mortgage debt, while additional information about other debt and asset holdings as well as self-assessed house value is available in the Wealth Supplement. We define liquid assets in as the total value of checking/savings accounts, bonds, and public equity holdings. Our definition of income is disposable income (wage income, transfers, business income, interest and dividend income, less taxes); we drop households with income less than \$1000/year. Mortgage debt is identified as outstanding balance of the first mortgage on the primary residence, and HELOC debt includes second mortgage and all other outstanding junior lien loan balances on the primary residence. We identify a refinancing event if the origination year of the first mortgage on the primary residence in a subsequent wave is greater than that recorded in the previous wave while the primary residence remains unchanged.

E Estimated model: inspecting the mechanism

E.1 Sensitivity analysis

Here analyze the sensitivity of the simulated moments to the estimated parameters, which underpins our structural identification. Table A.2 displays the values of simulated moments for different values of the key parameters in Θ, compared to the baseline case. For each of the seven estimated parameters we consider two values equidistant from the point estimates in either direction. Our discussion focuses on the key effect of each of the parameters.

Subjective discount factor δ Making households more patient via a larger δ increases the prevalence of homeownership, and increases household savings in the form of liquid asset holdings and home equity while lowering average mortgage balances). HELOC balances stay essentially the same (even though HELOC is more expensive than the mortgage on average in terms of the interest rate, it can be cheaper to access when liquidity is needed). As mortgage balances decline with higher δ, so does the frequency of refinancing and the sensitivity of refinancing to interest rates (β_R^{REFI} closer to 0). When the benefit of interest savings from refinancing is small, only those suffering from large income shocks find it worthwhile to pay the fixed costs of refinancing, as evidenced by the higher loan-to-income ratios and cash-out share for the new loans after refinancing. Moreover, under higher δ, while households cash-out more following negative aggregate income shocks (more negative β_Z), the consumption growth is still more affected by income shocks (larger β_Z^C), suggesting that households save the cashed-out home equity rather than consuming it. Finally, the average consumption/income ratio is higher with more patient households, again due to the fact that they have accumulated more savings via liquid assets and home equity.

Coefficient of relative risk aversion γ Increasing the risk aversion leads to more precautionary savings in the forms of liquid asset holdings and home equity (through both higher homeownership and lower mortgage balances), but also reduces the usage of HELOC as households accumulate enough liquid assets. Refinancing is mainly driven by the need to withdraw home equity rather than the purely financial incentive of lowering the mortgage rate, as cash-out/refi ratios increase in risk aversion and the sensitivity of refinancing to mortgage rate β_R^{REFI} moves closer to 0. Like the patient households, risk-averse households also cash-out more following negative aggregate consumption shocks (more negative β_Z) and shocks to mortgage rates (more negative β_R).

Intertemporal elasticity of substitution ψ A higher IES lowers liquid asset holdings, increases mortgage balances, and raises consumption volatility. This is due to the reason that households are less concerned with smoothing consumption over time, and the effects are qualitatively similar to those of a lower risk aversion. However, while a lower risk aversion coefficient reduces homeownership (which is driven by weaker precautionary savings motive), a higher IES raises homeownership. This is because the higher IES makes the refinancing option associated with owning a house more valuable, whereby households can better take advantage of house price appreciation and drop in interest rate.

The IES is also important for the dynamics of refinancing and cash-out. With a higher ψ, households are more willing to substitute consumption over time, therefore both cash-out and consumption are responding more to the changes in interest rates, as shown in a more negative β_R and a larger β_R^C.

Cost of refinancing ϕ_0, ϕ_1 Raising the quasi-fixed cost ϕ_0 of refinancing reduces the frequency of refinancing while increasing the new loan size and its cash-out component. Since costly refinancing makes mortgages effectively more expensive, average mortgage balances decline, as does homeownership. Its effect on the total leverage is partly offset by higher HELOC balances. Since lower mortgage balance reduces the risk in the household balance sheet, the precautionary holding of liquid assets is also lower. Raising the proportional cost parameter ϕ_1 has very similar effects. It might appear surprising that higher proportional refinancing cost increases the average new loan size and the cash-out share. This is driven by the composition effect: households are less likely to refinance for the purpose of lowering mortgage rates (β_R^{REFI} is -0.83 with high ϕ_1, compared to -1.09 in baseline case) but more likely to refinance to cash out home equity.

Table A.2: Sensitivity Analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
	Baseline	δ		γ		ψ		ν		ϕ_0		ϕ_1	
		0.874	0.966	1.518	6.071	0.151	0.603	0.067	0.268	0.077	0.309	0.007	0.028
All Households:													
Consumption/Income	0.71	0.68	0.74	0.69	0.74	0.71	0.72	0.73	0.70	0.72	0.71	0.71	0.71
Cons. growth vol, %	16.5	17.2	15.7	17.6	15.2	15.5	19.3	15.8	19.2	16.9	16.3	16.6	16.4
Homeownership rate, %	67.4	43.6	96.3	51.6	88.0	58.5	86.2	53.6	100.0	76.0	64.3	68.3	66.0
Homeowners:													
Liquid assets/Income	0.24	0.20	0.31	0.20	0.35	0.30	0.20	0.24	0.28	0.26	0.15	0.25	0.22
Mortgage/Income	0.94	1.34	0.56	1.34	0.51	0.92	1.07	0.91	1.43	1.22	0.45	1.00	0.84
HELOC/Income	0.09	0.08	0.09	0.08	0.07	0.07	0.11	0.08	0.09	0.08	0.12	0.08	0.09
Refinancing rate, %	11.1	17.4	6.6	17.2	6.2	11.5	12.5	11.0	15.9	17.2	4.8	12.0	9.7
Refi loan/Income	2.73	2.39	2.97	2.47	2.91	2.57	2.81	2.64	2.83	2.51	2.97	2.71	2.76
Dollar cash-out/Refi loan	0.51	0.42	0.56	0.41	0.60	0.52	0.49	0.51	0.46	0.43	0.62	0.50	0.54
Renters:													
Liquid assets/Income	0.15	0.07	0.38	0.08	0.34	0.18	0.19	0.13	0.09	0.16	0.14	0.15	0.15
Refinancing Regression:													
Coefficient on R, β_R^{REFI}	-0.97	-1.40	-0.66	-1.46	-0.56	-0.61	-1.54	-0.80	-1.35	-1.57	-0.26	-1.03	-0.83
Cash-out Regression:													
Coefficient on R, β_R	-0.57	-0.25	-0.84	-0.32	-0.71	-0.26	-1.14	-0.35	-1.08	-0.72	-0.24	-0.54	-0.59
Coefficient on Z, β_Z	-0.19	-0.09	-0.24	-0.09	-0.20	-0.12	-0.38	-0.25	0.13	-0.18	-0.38	-0.18	-0.25
Coefficient on H, β_H	0.10	0.14	0.01	0.15	0.04	0.06	0.17	0.10	0.05	0.08	0.06	0.11	0.08
Aggregate Consumption:													
Growth volatility, %	3.9	3.9	4.0	4.0	4.1	4.0	4.5	3.6	5.3	4.2	3.7	3.9	3.8
Sensitivity to Z, β_Z^C	1.30	1.25	1.37	1.28	1.36	1.34	1.40	1.22	1.72	1.41	1.26	1.30	1.28
Sensitivity to H, β_H^C	0.09	0.14	0.04	0.13	0.06	0.10	0.08	0.10	0.06	0.08	0.08	0.09	0.09
Sensitivity to lagged r, β_r^C	0.13	0.06	0.20	0.09	0.18	0.07	0.26	0.09	0.24	0.15	0.12	0.13	0.13
Sensitivity to lagged R, β_R^C	0.17	0.08	0.26	0.12	0.23	0.07	0.39	0.12	0.33	0.22	0.15	0.18	0.17